VMware Cloud Foundation on Dell EMC VxRail Admin Guide

VMware Cloud Foundation 3.9

You can find the most up-to-date technical documentation on the VMware website at:

https://docs.vmware.com/

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

About VMware Cloud Foundation on Dell EMC VxRail

1

The VMware Cloud Foundation on Dell EMC VxRail Administration Guide provides information on managing the integration of VMware Cloud Foundation and Dell EMC VxRail. As this product is an integration of VMware Cloud Foundation and Dell EMC VxRail, the expected results are obtained only when the configuration is done from both the products. This guide covers all the information regarding the VMware Cloud Foundation workflow. For the instructions on configuration to be done on Dell EMC VxRail, this guide provides links to the Dell EMC VxRail documentation.

Intended Audience

The *VMware Cloud Foundation on Dell EMC VxRail Administration Guide* is intended for the system administrators of the VxRail environments who want to adopt VMware Cloud Foundation. The information in this document is written for experienced data center system administrators who are familiar with:

- Concepts of virtualization, software-defined data centers, virtual infrastructure (VI), and virtual desktop infrastructure (VDI)
- VMware virtualization technologies, such as VMware ESXi™, the hypervisor
- Software-defined networking using VMware NSX®
- Software-defined storage using VMware vSAN™
- IP networks

Additionally, you should be familiar with these software products, software components, and their features:

- Dell EMC VxRail Manager
- VMware vSphere®
- VMware vCenter Server® and VMware vCenter Server® Appliance™
- VMware Platform Services Controller™
- VMware vRealize® Log Insight™
- vRealize Operations Manager
- vRealize Automation

Related Publications

The *VMware Cloud Foundation Planning and Preparation Guide* provides detailed information about the software, tools, and external services that are required for Cloud Foundation.

The *VMware Cloud Foundation Architecture and Deployment Guide* contains detailed information about a Cloud Foundation system, its components, and the network topology of a deployed system.

The *VMware Cloud Foundation Operations and Administration Guide* provides information about managing a VMware Cloud Foundation™ system, including managing the system's virtual infrastructure, managing users, configuring and deploying service offerings, and upgrading and monitoring the system.

Administering VMware Cloud Foundation on Dell EMC VxRail

2

VMware Cloud Foundation on Dell VMC VxRail enables VMware Cloud Foundation SDDC Manager on top of the Dell EMC VxRail platform.

An administrator of a VMware Cloud Foundation on Dell EMC VxRail system performs tasks such as:

- Manage certificates.
- Add capacity to your system.
- Configure and provision the systems and the workload domains that are used to provide service offerings.
- Manage provisioned workload domains.
- Monitor alerts and the health of the system.
- Troubleshoot issues and prevent problems across the physical and virtual infrastructure.
- Perform life cycle management on the software components.

Imaging of the Management Nodes 3

Image the management nodes by using Dell EMC RASR (Rapid Appliance Self Recovery) process. Ensure that you update the RASR image in each server node SD card before you start the imaging process.

For information on how to image the VxRail nodes, contact Dell EMC Support.

VxRail First Run for the Management Cluster

4

The VxRail first run for the management cluster consists of the following tasks:

- The discovery of the VxRail Nodes occurs. All the nodes that were imaged are detected.
- Upload the JSON configuration file. Trigger the validation.
- All the configuration inputs are validated.

The following components are deployed and enabled:

- vCenter
- VSAN
- VxRail Manager

Click **Manage VxRail** to log in to the VMware vCenter server.

For information on VxRail First Run, contact Dell EMC Support.

Change the Network Binding Type for the Management Network Port Group

5

Dell EMC VxRail configures the management network port with static binding which is the default setting. But as per the VMware Validated Design guidelines, you need to change the binding type to ephemeral for the management network port.

To change the binding type for the management network port:

Prerequisites

1 Log in to vCenter Server Appliance through vSphere Web client (Flash version).

2 Navigate to **Networking->Distributed Switch**.

3 Right click the **Management Network Port** group and click **Edit Settings**.

4 Verify if the network binding type is static binding and note down the VLAN ID.

Procedure

1 In vCenter Server Appliance, click **Networking**.

2 Select **Distributed Switch** > **Distributed Port Group**.

   Right-click **New Distributed Port Group**, update **Name** as **TempGroup** and keep the configuration same as the management network port group.

3 Change **VLAN type** from **None** to **VLAN**.

4 Update **VLAN ID** to that of the management network port group.

5 Click **Next** and review the configurations.

6 Click **Finish**.

   A new port group, **TempGroup**, is created and is shown in the port group list of the distributed switch.

7 Click **Hosts and Clusters**.

8 Select the first host of the cluster and click **Configure**.

9 Click **Virtual Switches > Distributed Switch**.

10 Click the **Migrate physical to virtual network adapters to this distributed switch** option.

11 In the **Migrate Networking** pop up screen, select only **Manage VMkernel adapters**.

Click **Next**.

12 Select the VMkernel network adapter associated with the management network.

13 Click **Assign port group**.

14 Select **TempGroupAssign destination port group** pop up screen and click **OK**.

Click **Next**.

15 Confirm **Overall impact** status as **No impact**. Click **Finish**.

16 Repeat steps 8 to 15 for all the other hosts of the cluster.

17 Click **Networking**.

18 Right-click **Management Network** and click **Edit Settings**.

In the **Edit Settings** screen, change **Port binding** from **Static binding** to **Ephemeral - no binding**.

19 Select **VLAN** and make sure **VLAN type** is **VLAN ID** is the same.

20 Click **Hosts and Clusters**.

Select the first host of the cluster and click **Configure**.

21 Select **Virtual Switches > Distributed Switch**.

22 Click the **Migrate physical to virtual network adapters to this distributed switch** option.

23 In the **Migrate Networking** pop up screen, select only **Manage VMkernel adapters**.

Click **Next**.

24 Select the VMkernel network adapter associated with the management network.

25 Click **Assign port group**.

26 Select the **Management Network in Assign destination port group** pop up screen and click **OK**. Click Next.

27 Confirm **Overall impact status** as **No impact**. Click **Finish**.

28 Repeat steps 21 to 27 for all the other hosts of the cluster.

29 Click **Networking**.

30 Select **Distributed Switch > TempGroup**.

31 Right-click **TempGroup** and click **Delete**.

32 Click **Yes** in **Delete Distributed Port Group** pop up UI screen.

Externalize vCenter Server

6

By default, VMware vCenter server is deployed within the VxRail cluster during the first run process. It must be externalized so that you can manage and lifecycle through SDDC Manager.

To convert the embedded VxRail VMware vCenter server to a customer-managed VMware vCenter server, contact Dell EMC Support.

Deploy VMware Cloud Builder Appliance 7

The Cloud Builder VM is a VM which also includes a service called the VMware Imaging Appliance which can be utilized for installing the base ESXi operating system on your physical servers. After you image the servers, you use the Cloud Builder to deploy and configure the management domain and transfer inventory and control to SDDC Manager. During the deployment process, the Cloud Builder validates network information you provide in the deployment parameter workbook such as DNS, network (VLANS, IPs, MTUs), and credentials.

You must deploy the Cloud Builder VM on a suitable platform. This can be on a laptop running VMware Workstation or VMware Fusion, or on an ESXi host. The Cloud Builder VM must have network access to all hosts on the management network.

The procedure here describes deploying the Cloud Builder VM on an ESXi host. Other deployment methods have different procedures.

Prerequisites

The Cloud Builder requires the following resources.

Component	Requirement
CPU	4 vCPUs
Memory	4 GB
Storage	150 GB

To automate the deployment, the Cloud Builder VM must be on the same management network as the hosts to be used. It must also be able to access all required external services, such as DNS and NTP.

Procedure

1 Download the Cloud Builder VM OVA on the Windows machine.

2 Log in to the vSphere Host Client.

3 In the navigator, select **Host**.

4 Click **Create/Register VM**.

5 On the Select creation type dialog box, select **Deploy a virtual machine from an OVF or OVA file** and click **Next**.

6. Enter a name for the VM.
7. Select **Click to select files or drag/drop**. Select the Cloud Builder VM OVA from your local file system and click **Open**.
8. Click **Next**.
9. On the Select Storage page, select the storage for the Cloud Builder VM.
10. On the License agreements dialog box, click **I agree** and then click **Next**.
11. On the Select networks dialog box, select the port group associated with the VLAN ID used by the ESXi hosts where Cloud Foundation will be deployed and then click **Next**.
12. On the Additional settings dialog box, expand **Application** and enter the following information for the Cloud Builder VM:

Setting	Details
Admin Username	The admin user name cannot be one of the following pre-defined user names: ■ root ■ bin ■ daemon ■ messagebus ■ systemd-bus-proxy ■ systemd-journal-gateway ■ systemd-journal-remote ■ systemd-journal-upload ■ systemd-network ■ systemd-resolve ■ systemd-timesync ■ nobody ■ sshd ■ named ■ rpc ■ tftp ■ ntp ■ smmsp ■ cassandra
Admin Password/Admin Password confirm	The admin password must be a minimum of 8 characters and include at least one uppercase, one lowercase, one digit, and one special character.
Root password/Root password confirm	The root password must be a minimum of 8 characters and include at least one uppercase, one lowercase, one digit, and one special character.
Hostname	Enter the hostname for the Cloud Builder VM.
Network 1 IP Address	Enter the IP address for the Cloud Builder VM.
Network 1 Subnet Mask	For example, `255.255.255.0`.
Default Gateway	Enter the default gateway for the Cloud Builder VM.

Setting	Details
DNS Servers	IP address of the primary and secondary DNS servers (comma separated). Do not specify more than two servers.
DNS Domain Name	For example, `vsphere.local`.
DNS Domain Search Paths	Comma separated. For example `vsphere.local, sf.vsphere.local`.
NTP Servers	Comma separated.

13 Review the deployment details and click **Finish**.

Note Make sure your passwords meet the requirements specified above before clicking **Finish** or your deployment will not succeed.

14 After the Cloud Builder VM is deployed, SSH in to the VM with the admin credentials provided in step 12.

15 Ensure that you can ping the ESXi hosts.

16 Verify that the Cloud Builder VM has access to the required external services, such as DNS and NTP by performing forward and reverse DNS lookups for each host and the specified NTP servers.

Initiate the Cloud Foundation Bring-Up Process

8

During bring-up, the management domain is created on the ESXi hosts specified in the deployment configuration spreadsheet. The Cloud Foundation software components are automatically deployed, configured, and licensed using the information provided.

Procedure

1. Log in to the Cloud Builder VM with your admin credentials by navigating to `https://`*`Cloud_Builder_VM_IP`*.
2. Verify that the criteria mentioned in the bring-up checklist have been met. Select **Check All** and click **Next**.
3. Read the End-User License Agreement and accept it. Click **Next**.
4. You can download the deployment parameter sheet. The deployment parameter spreadsheet provides a mechanism to specify the required deployment information specific to your environment. This includes information about your networks, hosts, license keys, and other information. The spreadsheet is downloaded from the Cloud Builder VM. The completed spreadsheet is then converted to a JSON file. The deployment parameter spreadsheet can be reused to deploy multiple Cloud Foundation instances of the same version. Click **Download Deployment Parameter Sheet**. Complete the worksheet.

   Note Ensure that you enter a complete cluster name or the corresponding prefix as in the vCenter Server. Regular expression is not supported for the cluster name parameter.

5. To upload the config file, click **Upload** . You can choose the `.xlsx` or `.json` file for updating the configuration. Ensure that all the IP addresses are reserved in the DNS. Select the completed parameter sheet and click **Upload**.
6. After the file is uploaded, click **Validate** to validate the uploaded file.

   The following audit tasks are performed and validation results are displayed on the UI.

   - JSON specifications validation: Validates the completeness and correctness of the specifications of JSON.
   - Well-Formed JSON File: Validates JSON correctness, syntax, null values, and missing fields or components.
   - Password validation

Validates specified passwords. Checks for minimum length, invalid characters, and format.

- ESXi host version validation

  Validates ESXi version installed on the hosts and compares against the VCF-EMS manifest located in `/opt/evosddc/bundle/scripts/manifest.json` on the Cloud Foundation Builder VM.

- Cloud Builder Readiness: Validates whether the requirements to run the Cloud Foundation Builder VM are met.

- License key format

  Validates format, validity, and expiry for ESX, vSAN, vCenter Server, NSX, and Log Insight license keys.

- ESXi Host Readiness

- Network configuration: Validates CIDR to IP address validity, IP addresses in use, gateways, invalid or missing VLANs, invalid or missing MTU, and network spec availability for all components.

- Time Synchronization: Validates the time on the components is synchronized with the NTP server in the SDDC Manager.

- Host and IP DNS records: Validates SSH access and policy, NTP configuration and policy, DNS configuration, VMNIC availability, vSwitch availability, VM network portgroup , and VLAN check on each host.

To access the bring-up log file, SSH to the Cloud Builder VM as root and open the `/opt/vmware/bringup/logs/vcf-bringup-debug.log` file.

If there is an error during the validation and the **Next** button is grayed out, you can either make corrections to the environment or edit the JSON file and upload it again. Then click **Re-Try** to perform the validation again.

If any warnings are displayed and you want to proceed, click **Acknowledge** and then click **Next**.

7 Click **Begin Bring-Up**.

During the bring-up process, the following tasks are completed.

- The second PSC controller, SDDC Manager, NSX components, and vRealize Log Insight are deployed.

- The management domain is created, which contains the management components - SDDC Manager, vCenter Server, and NSX Managers and Controllers.

- The VxRail Manager version is fetched.

The status of the bring-up tasks is displayed in the UI. You can download the list of tasks by clicking **Download**.

After bring-up is completed, a green bar is displayed indicating that bring-up was successful. A link to the SDDC Manager UI is also displayed.

If there are errors during the bring-up, see Troubleshooting Cloud Foundation Deployment for guidance on how to proceed.

8 Navigate to the SDDC Manager Dashboard by clicking the link displayed on the UI.

9 Verify the following:

- View the management domain details.
- Log in to vCenter Server and verify the management cluster, vSAN cluster, and deployed VMs.

10 Power off the Cloud Builder VM.

If you no longer need the Cloud Foundation Builder VM, you can delete it. If you delete it now, you can always redeploy it later.

Manage Certificates

9

You can change the certificates of all VxRail Manager instances as well as the certificates for all Cloud Foundation components through SDDC Manager. See Managing Certificates for Cloud Foundation Components in *VMware Cloud Foundation Operations and Administration Guide*.

Manage Passwords

10

You specify the passwords for your Cloud Foundation system's internal accounts as part of the bring-up procedure. You can also modify the passwords for these accounts using RESTful API calls.

You can update or rotate the password for the `root` and `mystic` users of the VxRail Manager and the `root` user of ESXi hosts using the SDDC Manager. To update or rotate the passwords for other users refer to the Dell EMC VxRail documentation.

To provide the optimal security and proactively prevent any passwords from expiring, you should rotate passwords every 80 days.

This chapter includes the following topics:

Configure Dual Authentication

You must configure dual authentication in order to perform certain tasks, such as updating or rotating passwords and configuring NSX Manager backups.

You will use the vSphere Client to create a new SSO group (`Sddc_Secured_Access`), add a user to the group, and assign a password to that user. The user is called the privileged user and will be required, along with its password, to perform certain tasks from the SDDC Manager UI or the VMware Cloud Foundation API.

You can create a new SSO user as the privileged user, or use an existing SSO user. If you plan to invoke operations requiring the privileged user as part of an automation solution, you should create a separate SSO user for this purpose. The SSO users used by automation should also be assigned the `No Access` role.

Note The **`administrator@vsphere.local`** user cannot be the privileged user.

Prerequisites

To perform this operation, you need to log in to the management vCenter Server as the `administrator@vsphere.local` user or another user who has the administrator role.

Procedure

1. Log into management vCenter Server using the vSphere Client.
2. Navigate to **Administration > Single Sign On > Users and Groups**.
3. Click the **Users** tab and select the domain from the drop-down list.
4. To create a new user in the selected domain, click **Add User**, enter the required information, and click **Add**.
5. Click the **Groups** tab and click **Add Group**.
6. Create a group named `Sddc_Secured_Access`, add the new or existing user to the group, and click **Add**.

Rotate Passwords for Managed Entities

As a security measure, you can rotate passwords for the logical and physical entities on all racks in your system. The process of password rotation generates randomized passwords for the selected accounts.

You can rotate passwords for the following entities.

- ESXi
- vCenter Server

  By default, the vCenter Server root password expires after 90 days.
- PSC

  By default, the PSC password expires after 90 days.
- NSX Manager
- NSX Controllers (NSX for vSphere and NSX-T)
- NSX Edge
- NSX-T Manager
- vRealize Log Insight
- vRealize Operations
- vRealize Automation
- vRealize Suite Lifecycle Manager
- SDDC Manager **`backup`** user

To update the SDDC Manager root, super user, and API passwords, see Updating SDDC Manager Passwords.

Prerequisites

- Verify that there are no currently failed workflows in your Cloud Foundation system. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no active workflows are running or are scheduled to run during the brief time period that the password rotation process is running. It is recommended that you schedule password rotation for a time when you expect to have no running workflows.
- Configure the privileged user. For more information, see Configure Dual Authentication.

Procedure

1. From the navigation pane, choose **Administration > Security > Password Management > Locally Managed.**.

   The Password Management page displays a table with detailed information about all domains, including their component, credential type, FQDN, IP address, and user name. This table is dynamic. Each column can be sorted.

   You can click the filter icon next to the table header and filter the results by a string value. For example, click this icon next to **User Name** and enter `admin` to display only domains with that user name value.

2. Select the component type for which you want to rotate passwords from the **Component** drop-down menu. For example, **ESXI**.
3. Select one or more components and click **Rotate**.
4. Enter the privileged username and the privileged password and click **Rotate**.

   A message appears at the top of the page showing the progress of the operation. The Tasks panel also shows detailed status for the password rotation operation. Click on the task name to view sub-tasks. As each of these tasks are run, the status is updated. If the task fails, you can click **Retry**.

Results

Password rotation is compete when all sub-tasks are completed successfully.

Manually Update Passwords

You can manually change the password for a selected domain account. Unlike password rotation, which generates a randomized password, you provide the new password.

You can modify only one password at a time.

Note You cannot use these controls to update the NSX-T password. You can only update the NSX-T password from the NSX-T Manager product interface.

Prerequisites

- Verify that there are no currently failed workflows in your Cloud Foundation system. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no active workflows are running or are scheduled to run during the manual password update.
- Configure the privileged user. For more information, see Configure Dual Authentication.

Procedure

1 From the navigation pane, choose **Administration > Security > Password Management**.

   The Password Management page displays a table with detailed information about all domains, including their component, credential type, FQDN, IP address, and user name. This table is dynamic. Each column can be sorted.

   You can click the filter icon next to the table header and filter the results by a string value. For example, click this icon next to **User Name** and enter `admin` to display only domains with that user name value.

2 Select the domain entity whose password you want to update and click **Update** at the top of the page.

   Note If you select more than one domain, the **Update** button is disabled.

   The Update Password dialog box appears. This dialog box also displays the entity name, credential type, user name, privileged user name, privileged password in case you need to confirm you have selected the correct domain. Enter the values for all these fields.

3 Enter the privileged username and privileged password.

4 Enter and confirm the new password.

   If the passwords, do not match, the dialog displays a red alert.

5 Click **Update**.

   A message appears at the top of the page showing the progress of the operation. The Tasks panel also shows detailed status of the password update operation. Click on the task name to view sub-tasks.

   If the Tasks panel shows the task as having failed, click **Retry**.

Results

Password update is compete when all sub-tasks are completed successfully.

Look Up Account Credentials

To look up the account credentials for the built-in accounts that are managed and rotated by SDDC Manager, you log in to the SDDC Manager VM using the root account credentials.

Prerequisites

You must have the root account credentials to log in to the SDDC Manager VM.

Configure the privileged user. For more information, see Configure Dual Authentication.

Procedure

1. SSH in to the SDDC Manager VM using the **vcf** user account.
2. (Optional) Change to the /usr/bin directory.

   Note Although the password management CLI commands are located in /usr/bin, you can run them from any directory.
3. Enter su to switch to the root user.
4. Obtain the account credentials list by typing the command:

   ```
   lookup_passwords
   ```

   You will be required to enter the privileged user name and the privileged password.

   To display the output in JSON format, use the following example command:

   ```
   curl “https://localhost/security/password/vault” -k -u "<administrative user
   name>:<password>" -H “Accept: application/json” -H “privileged-username: vcf-secure-
   user@vsphere.local” -H “privileged-password: AfGh!8f9”
   ```

   Enter the required credentials.
5. (Optional) Save the command output to a secure location with encryption so that you can access it later and use it to log in to the components as needed.

Password Management cURL API Reference

You can perform basic password management operations using cURL API requests. SSH in to the SDDC Manager VM and log in as the root user to use the cURL API.

cURL Password Operation API Requests

Some of the above operations can be run using cURL API requests.

Look up passwords - JSON format

Retrieves and lists in JSON format the account credentials for the built-in accounts that are managed and rotated by SDDC Manager.

```
# curl 'http://localhost/security/password/vault' \
     -i -H 'Accept: application/json'
```

Look up passwords - plain text format

Retrieves and lists in plain text format the account credentials for the built-in accounts that are managed and rotated by SDDC Manager.

```
# curl 'http://localhost/security/password/vault' \
      -i  -H 'Accept: text/plain'
```

Update password

Updates the password of the specified domain component.

```
# curl 'http://localhost/security/password/vault' -i -X POST \
     -H 'Content-Type: application/json' \
     -H 'Accept: application/json' \
     -d '{
   "entities": [{
       "credentialType" : "<credential type such SSH or API>",
       "entityIpAddress" : "<IP address>",
       "entityType" : "<component, such as ESXI>",
       "entityId" : "<node ID value>",
       "password" : "<password>",
       "domainName" : "<domain name>",
       "entityName" : "<FQDN>",
       "username" : "root"
    }],
   "type":"UPDATE"
}'
```

Rotate password

Rotates the password of the specified domain component.

```
# curl 'http://localhost/security/password/vault' -i -X POST \
     -H 'Content-Type: application/json' \
     -H 'Accept: application/json' \
     -d '{
   "entities": [{
       "credentialType" : "<credential type such SSH or API>",
       "entityIpAddress" : "<IP address>",
       "entityType" : "<component, such as ESXI>",
       "entityId" : "<node ID value>",
       "password" : "<password>",
       "domainName" : "<domain name>",
       "entityName" : "<FQDN>",
       "username" : "root"
    }],
   "type":"ROTATE"
}'
```

Password operation history

Returns in JSON format the password history recorded in the password management database.

```
# curl 'https://localhost/security/password/vault/transactions' \
     -i -H 'Accept: application/json' \
     -k -u "<administrative user name>:<password>"
```

Password operation status

Returns in JSON format the latest (or current) workflow, which is an asynchronous job running in SDDC Manager. It polls the status of the workflow and reports percentage completed until the workflow finishes, at which time it reports its status.

```
# curl 'https://localhost/security/password/vault/transactions/2002' \
     -i -H 'Accept: application/json'\
     -k -u "<administrative user name>:<password>"
```

Retry failed password operation

Retries the specified failed operation and returns results in JSON format

```
# curl 'http://localhost/security/password/vault/transactions/2002' \
     -i -X PATCH \
     -H 'Content-Type: application/json' \
     -H 'Accept: application/json' \
     -d '{
    "entities": [{
        "credentialType" : "<credential type such SSH or API>",
        "entityIpAddress" : "<IP address>",
        "entityType" : "<component, such as ESXI>",
        "entityId" : "<node ID value>",
        "password" : "<password>",
        "domainName" : "<domain name>",
        "entityName" : "<FQDN>",
        "username" : "root"
     }],
    "type":"<specify ROTATE or UPDATE>"
}'
```

Cancel password operation

Cancels failed password operations and returns results in JSON format

```
# curl 'https://localhost/security/password/vault/transactions/2002' \
     -i -X DELETE -H 'Accept: application/json' \
     -k -u "<administrative user name>:<password>"
```

Updating SDDC Manager Passwords

You cannot update SDDC Manager passwords through the SDDC Manager Dashboard or by using cURL API requests. Instead, you will need to SSH into the SDDC Manager VM and make the changes there.

The process for updating SDDC Manager passwords varies, depending on which account you are updating.

- Update SDDC Manager Root and Super User Passwords

  For security reasons, you can change passwords for the SDDC Manager root (**root**) and super user (**vcf**) accounts. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

- Update SDDC Manager REST API Account Password

  To use the VMware Cloud Foundation API, an API client logs in using the SDDC Manager **admin** account. For security reasons, you should periodically update the password for this account.

Update SDDC Manager Root and Super User Passwords

For security reasons, you can change passwords for the SDDC Manager root (**root**) and super user (**vcf**) accounts. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

The SDDC Manager **root** password expires after 365 days.

Procedure

1 SSH in to the SDDC Manager VM using the **vcf** user account.

2 Enter **su** to switch to the root user.

3 Enter one of the following commands:

Option	Description
`passwd vcf`	To change the super user password.
`passwd root`	To change the root password.

4 Enter and retype the new password. For example:

```
root@sddc-manager [ /home/vcf ]# passwd vcf
New password:
Retype new password:
passwd: password updated successfully
```

Results

The password is updated.

Update SDDC Manager REST API Account Password

To use the VMware Cloud Foundation API, an API client logs in using the SDDC Manager **admin** account. For security reasons, you should periodically update the password for this account.

If you write a script that invokes the APIs, the script should either prompt the user for the password for the **admin** account or should accept the password as a command line option. As a best practice, you should not encode the password for the account in the script code itself.

Password requirements:

- Length 8-12 characters
- Must include: mix of upper-case and lower-case letters a number a special character such as @ ! # $ % ^ or ?
- Cannot include: * { } [ ] ( ) / \ ' " ` ~ , ; : . < >

Procedure

1 SSH in to the SDDC Manager VM using the **vcf** user account.

2 Enter **su** to switch to the root user.

3 Enter the following command:

```
/opt/vmware/vcf/commonsvcs/scripts/auth/set-basicauth-password.sh admin <password>
```

For *<password>*, enter the new password for the **admin** account.

About VI Workload Domains

11

In the VI Configuration wizard, you specify the storage, name, compute, and NSX platform details for the VI workload domain. Based on the selected storage, you provide vSAN parameters or NFS share details. You then select the hosts and licenses for the workload domain and start the creation workflow.

The workflow automatically:

- Deploys an extra vCenter Server Appliance for the new workload domain within the management domain.

  By using a separate vCenter Server instance per workload domain, software updates can be applied without impacting other workload domains. It also allows for each workload domain to have additional isolation as needed.

- Connects the specified ESXi servers to this vCenter Server instance and groups them into a cluster. Each host is configured with the port groups applicable for the workload domain.

- Configures networking on each host.

- Configures vSAN or NFS storage on the ESXi hosts.

- Deploys an NSX Manager in the management domain and three NSX controllers on the ESXi datastore for each NSX for vSphere workload domain . The workflow also configures an anti-affinity rule between the controller VMs to prevent them from being on the same host for High Availability.

- Deploys a cluster of three NSX-T Managers in the management domain for the first NSX-T VI workload domain. The workflow also configures an anti-affinity rule between the NSX Manager VMs to prevent them from being on the same host for High Availability. All subsequent NSX-T workload domains share this NSX-T Manager cluster.

  For an NSX-T workload domain, NSX Edges are needed to enable overlay VI networks and public networks for the north-south traffic. NSX Edges are not deployed automatically for an NSX-T VI workload domain. You can deploy them manually after the VI workload domain is created. Subsequent NSX-T VI workload domains share the NSX-T Edges deployed for the first workload domain.

- NSX Managers deployed as part of a VI workload domain are configured to periodically get backed up to an SFTP server. By default, these backups are written to an SFTP server built into SDDC Manager, but you can register an external SFTP server for better protection against failures. See "Configure an External SFTP Server for File-Based Backups" in the *VMware Cloud Foundation Operations and Administration Guide*. SDDC Manager uses either the built-in or external SFTP server with all currently deployed NSX Managers and when deploying additional NSX Managers.
- Licenses and integrates the deployed components with the appropriate pieces in the Cloud Foundation software stack.

Note Currently, the NSX-V to NSX-T migration is not supported in the VMware Cloud Foundation environments.

This chapter includes the following topics:

Prerequisites for a VI Workload Domain

This section lists pre-requisites for a VI workload domain.

- A DHCP server must be configured on the NSX-T host overlay (Host TEP) VLAN of the workload domain. When NSX-T creates Edge Tunnel End Points (TEPs) for the VI workload domain, they are assigned IP addresses from the DHCP server.
- A minimum of three hosts must be available for the cluster to be created during the VxRail first run.
- Decide on a name for your VI workload domain. Each VI workload domain must have a unique name. It is good practice to include region and site information in the name because resource object names (such as host and vCenter names) are generated based on the VI workload domain name. The name can be three to 20 characters long and can contain any combination of the following:
  - Lowercase alphabetic characters
  - Uppercase alphabetic characters
  - Numbers

  Note Spaces are not allowed in any of the names you specify when creating a VI workload domain.

- Decide on the following passwords - refer to the appropriate table for the Cloud Foundation version in your environment.
  - vCenter root password
  - NSX-T Manager admin password
- Update the management domain before you deploy `WLD`. To check the management domain Bill of Materials (BOM), see *VMware Cloud Foundation on Dell EMC VxRail Release Notes*. If the management domain BOM does not match, connect to the LCM depot, download, and apply the upgrade patches.
- Table 11-1. Passwords for Cloud Foundation

Account	Password Requirements
vCenter root	1 Length 8-20 characters 2 Must include: ■ mix of upper-case and lower-case letters ■ a number ■ a special character
NSX-T Manager admin	1 Minimum length 12 characters 2 Must include: ■ at least one lowercase and one uppercase letter ■ a number ■ a special character ■ exclude_char such as `{ } [ ] ( ) / \ ' " ` ~ , ; : . < >` ■ at least five different characters 3 Must not include: ■ a dictionary word ■ a palindrome ■ more than four monotonic character sequences

- Gather the information that you need for the workload domain creation workflow.

vCenter IP address and FQDN
Three NSX Managers IP addresses and FQDNs
NSX Manager Virtual IP (VIP) address and FQDN

- The IP addresses and Fully Qualified Domain Names (FQDNs) for the vCenter and NSX Manager instances must be resolvable by DNS.
- You must have valid license keys for the following products:
  - NSX-T Data Center
  - vSAN

    Because vSAN licenses are per CPU, ensure that you have sufficient licenses for the ESXi hosts to be used for the workload domain.

Create a VxRail Virtual Infrastructure Workload Domain

You have to create a VxRail Virtual Infrastructure (VI) workload domain before adding a cluster.

Deploy the vCenter server and make the domain ready for the cluster addition.

Note You can only perform one workload domain operation at a time. For example, while creating a new workload domain, you cannot add a cluster to any other workload domain.

For each NSX for vSphere workload domain, the workflow deploys an NSX Manager in the management domain and three NSX controllers on the ESXi datastore. The workflow also configures an anti-affinity rule between the controller VMs to prevent them from being on the same host for High Availability.

For the first NSX-T VI workload domain in your Cloud Foundation environment, the workflow deploys a cluster of three NSX-T Managers in the management domain. The workflow also configures an anti-affinity rule between the NSX-T Manager VMs to prevent them from being on the same host for High Availability. All subsequent NSX-T workload domains share this NSX-T Manager cluster.

For an NSX-T workload domain, the NSX-T edges are needed to enable overlay the VI networks and the public networks for north-south traffic. The NSX-T edges are not deployed automatically for an NSX-T VI workload domain. You can deploy them manually after the VI workload domain is created. The subsequent NSX-T VI workload domains may share the NSX-T Edges deployed for the first workload domain.

Start the VI Configuration wizard.

Procedure

- On the **SDDC Manager** Dashboard, click **+ Workload Domain** and then select **VI-VxRail Virtual Infrastructure Setup**.

Specify Name

Provide a name for the VxRail VI workload domain and organization.

Procedure

1. Type a name for the VI workload domain, such as `sfo01`. The name must contain between 3 and 20 characters.

   It is a good practice to include location information in the name as resource object names (such as host and vCenter names) are generated on the basis of the VI workload domain name.

2. Type a name for the organization that requested or will use the virtual infrastructure, such as `Finance`. The name must contain between 3 and 20 characters.

3. Click **Next**.

Specify Compute Details

Specify the compute (vCenter) details for this workload domain.

Procedure

1 On the Compute page of the wizard, enter the vCenter IP address and DNS name.

   Note Before updating the IP address in the wizard, ensure that you have reserved the IP addresses in the DNS.

2 Type the vCenter subnet mask and default gateway.

3 Type and re-type the vCenter Root password.

4 Click **Next**.

Review the Details

At the review step of the wizard, review the information about the workload domain and start the creation workflow. You can also print the information or download a printable version to print later.

The **Review** page displays information about the resources and their configurations that will be deployed when the workflow creates and deploys the virtual infrastructure for this workload domain.

Procedure

1 Scroll down the page to review the information.

2 Click **Finish** to begin the creation process.

   The **Workload Domains** page appears and a notification is displayed letting you know that VI workload domain is being added. Click **View Task Status** to view the domain creation tasks and sub tasks.

   If a task fails, you can fix the issue and re-run the task. If the workload domain creation fails, contact VMware Support.

Results

The status will be activating until we add the primary cluster in to domain. When the VxRail VI workload domain is created, it is added to the workload domains table along with the already listed management domain.

The OEM license is assigned by default to the workload domains.

Add the Primary VxRail Cluster

The creation of primary cluster is a part of the creation of the workload domain.

Before adding the VxRail cluster, you need to perform the imaging of the workload domain nodes. Once you complete the imaging, perform the VxRail first run of the workload domain nodes using the external vCenter Server.

- Create a local user in vCenter server as this is an external server deployed by VMware Cloud Foundation. This is required for the VxRail first run.

  a Log in to the workload domain vCenter Server Appliance through vSphere Web Client.

  b Select **Menu > Administration > Single Sign On**.

  c Click **Users and Groups**.

  d Click **Users**.

  e Select **Domain vSphere.local**.

  f Click **Add User**.

  g In the **Add User** pop-up window, enter the values for the mandatory fields.

  h Enter **Username** as `vxadmin` and **Password.** Confirm the **Password**.

  i Click **Add**.

  j Wait for the task to complete.

- Image the workload domain nodes. For information on imaging the nodes, contact Dell EMC Support.
- Do a VxRail first run of the workload domain nodes using the external vCenter Server. For information on the VxRail first run, contact Dell EMC Support.
- Once the validation is complete, trigger the build VxRail operation.
- The cluster is created in VxRail.

To add the primary VxRail cluster to a workload domain:

Procedure

1 On the **SDDC Manager** Dashboard, click **Inventory > Workload Domains**. The **Workload Domains** page displays information for all workload domains.

2 In the workload domains table, hover your mouse over workload domain in the activating state. The primary cluster needs to be added to the activating domain. This means that the domain is not created and it is waiting for the addition of primary cluster.

  A set of three dots appears on the left of the workload domain name.

3 Click these three dots. Click **Add VxRail Cluster**.

4 The **Add VxRail Cluster to Workload Domain** page appears.

5 On the **Discovered Clusters** page, a single VxRail cluster or multiple VxRail clusters in the vCenter are discovered. If there are multiple clusters, select a cluster. Click **Next**.

6 The **Discovered Hosts** page displays a list of the discovered hosts for that cluster. Update the SSH password for the discovered hosts for that cluster. Click the icon next to the name of the host to get more information about it. Click **Next**.

7 On the **VxRail Manager** page, enter the `admin` and `root` user names and passwords.

8 The Networking page displays all the networking details for the cluster.

   a Select the NSX Platform. Provide the corresponding **VLAN ID**.

   b If you have selected NSX-V, provide the manager and the controller information.

   c If you have selected NSX-T, provide the managers' details and the cluster/virtual IP address.

   d Click **Next**.

   e Enter the license keys for NSX for vSphere and VMware vSAN. Click **Next**.

   f Review the details. Click **Finish**. The process of adding the VxRail cluster is triggered.

Deploy NSX-T Edge Cluster on VxRail

For an NSX-T VI workload domain, the NSX-T edges are required to enable overlay for the VI and the public networks for north-south traffic. Use the following procedure to deploy for the Cloud Foundation version in your environment.

Prerequisites

Ensure that you can access the following documents:

- NSX-T Data Center Installation Guide
- VVD Guidance for Deployment of VMware NSX-T for Workload Domains

Procedure

1 Create the following resource pools for the shared edge and the compute cluster.

   - NSX-T Edge devices that control the network traffic in and out of the workload domain

     Note In this section, we have used `sfo01-w02rp-sddc-edge` as an example.

   - NSX-T Edge devices that provide networking services to the tenant workloads in the workload domain

     Note In this section, we have used `sfo01-w02rp-user-edge` as an example.

- Tenant workloads in the workload domain

Note In this section, we have used `sfo01-w02rp-user-vm` as an example.

a Right-click the specific shared cluster and select **New Resource Pool**.

b In the New Resource Pool dialog box, enter the values for the particular edge resource pool and click OK. For the following table, `sfo01-w02rp-sddc-edge`, `sfo01-w02rp-user-edge`, and `sfo01-w02rp-user-vm` are used as examples.

Setting	Resource Pool 1	Resource Pool 2	Resource Pool 3
Name	sfo01-w02rp-sddc-edge	sfo01-w02rp-user-edge	sfo01-w02rp-user-vm
CPU-Shares	High	Normal	Normal
CPU-Reservation	N/A	N/A	N/A
CPU-Reservation Type	Expandable Selected	Expandable Selected	Expandable Selected
CPU-Limit	Unlimited	Unlimited	Unlimited
Memory-Shares	Normal	Normal	Normal
Memory-Reservation	32 GB	0	0
Memory Reservation Type	Expandable Selected	Expandable Selected	Expandable Selected
Memory-Limit	Unlimited	Unlimited	Unlimited

c Repeat the step to add the remaining resource pools.

2 Create Transport Zones for Uplink Traffic. Follow the VMware Validated Designs guidance to create two uplink transport zones with names that match your environment with same settings as shown in the following link. `sfo01-w-uplink01` and `sfo01-w-uplink02` are examples of the uplink transport zones.

https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-3D614E16-9A84-4C47-83B9-8E0233DD5FCD.html

3 Create Uplink Profiles. Use the following link to create two additional edge uplink profiles and an edge overlay profile. The uplink profile names and VLANs should match that of the environment. There is no need to create the ESXi uplink profile as this has already been created by the VMware Cloud Foundation automation with the name starting with host-uplink.

https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-75DDCAAA-85B8-40AB-A923-2A1C9A74FDAD.html

Note From the section in the above mentioned link

- Omit Step 5 in the following link as the two teaming policies have already been created by the VMware Cloud Foundaion automation.
- For step 6, use `uplink-1` and `uplink-2` to match the teaming policy names in the host-uplink profile.
- Skip step 7 as NIOC profile has already been created by VMware Cloud Foundation.

4 Create NSX-T Segments for Edge Traffic. Follow the VMware Validated Designs documentation to create the logical segments for the edge uplink traffic. In the example, these are `sfo01-w-nvds01-uplink0`, `sfo01-w-nvds01-uplink0`, `sfo01-w-uplink01`, `sfo01-w-uplink01`, and `sfo01-w-overlay`. The uplink VLANs should match that of the environment.

Note The `sfo01-esxi-vlan` transport zone corresponds to the `vlan-tz` deployed by VMware Cloud Foundation.

https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-5EA93314-8E9B-4FC7-9275-E3996426C246.html

5 Create an NSX-T Edge cluster profile with a name that matches the environment using the settings as per the VMware Validated Designs documentation as shown in the following link:

https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-6A1A706F-31B4-4A48-8DCA-E9AADC14EB2A.html

6 Deploy the NSX-T Edge Appliances. Ensure that you use the VxRail vDS Management Network Port group for the first network adapter (Source Network 0) and the net-overlay segment for the second adapter (Source Network 1). For the network adapters 2 and 3, use the trunked uplink segments created in step 3. For more information, see:

https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-FDEDA95C-CD1F-42D1-BB05-874EF65AECFB.html?hWord=N4IghgNiBcIE4FMDOB7ArnAxggBABxRSgF8g

7 Join the NSX-T Edge Nodes to the Management Plane. Follow VVD guidance from the following link:

https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-600EF430-BC56-49F0-9D6D-CD5788463947.html

8 Create an Anti-Affinity Rule for the NSX-T Edge Nodes in the Shared Edge and Compute Cluster. Follow VVD guidance from the following link:

   https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-12EC5F9D-254F-478A-9B24-1E05822F823D.html

9 Add the NSX-T Edge Nodes to the Transport Zones. Follow the VMware Validated Designs procedure below. Use theeEdge switch names from the drop down menu to populate the name fields. The overlay edge switch name starts with the `nvds` string.

   https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-72B0926E-0682-4F11-B85D-5D15974D074C.html

10 Create an NSX-T edge cluster. Follow the VMware Validated Designs guidance from the following link:

   https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-14D05472-90E0-4DB2-8069-F0B565D3824F.html

11 Create and configure the Tier-0 Gateway. Follow the VMware Validated Designs guidance from the following link:

   https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-8BCF3424-1157-4C64-A4B6-03886EC31E15.html

12 Create and configure Tier-1 Gateway. Follow the VMware Validated Designs guidance from the following link:

   https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-F03C18AC-B552-4922-9F5E-6AB63B4411C2.html

13 Verify BGP Peering and Route Redistribution. Follow the VMware Validated Designs guidance from the following link:

   https://docs.vmware.com/en/VMware-Validated-Design/5.0.1/com.vmware.vvd.sddc-nsxt-domain-deploy.doc/GUID-0136A321-BCF4-4740-AE56-A1F4CD14F6D4.html

Expand a Workload Domain

After you add the primary cluster, you can add more clusters to expand the workload domain.

Before adding the VxRail cluster, you need to perform the imaging of the workload domain nodes. Once you complete the imaging, perform the VxRail first run of the workload domain nodes using the vCenter Server for the workload domain.

- Create a local user in vCenter server as this is an external server deployed by VCF. This is required for the VxRail first run.
- Image the workload domain nodes. For information on imaging the nodes, refer to Dell EMC VxRail documentation.

- Do a VxRail first run of the workload domain nodes using the vCenter server for that workload domain. For information on VxRail first run, refer to the Dell EMC VxRail documentation.
- Once the validation is complete, trigger the build VxRail operation.

Add the VxRail Cluster

To add the VxRail cluster to a workload domain:

Procedure

1. On the **SDDC Manager** Dashboard, click **Inventory > Workload Domains**. The **Workload Domains** page displays information for all workload domains.
2. In the workload domains table, hover your mouse in the VxRail workload domain row.

   A set of three dots appears on the left of the workload domain name.
3. Click these three dots. Click **Add VxRail Cluster**.
4. The **Add VxRail Cluster to Workload Domain** page appears.
5. On the **Discovered Clusters** page, a single VxRail cluster or multiple VxRail clusters in the vCenter are discovered. If there are multiple clusters, select a cluster. Click **Next**.
6. The **Discovered Hosts** page displays a list of the discovered hosts for that cluster. Update the SSH password for the discovered hosts for that cluster. Click the icon next to the name of the host to get more information about it. Click **Next**.
7. On the **VxRail Manager** page, enter the `admin` and `root` user names and passwords.
8. The Networking page displays all the networking details for the cluster.
   a. Select the NSX platform. If you select NSX-T, ensure that you have downloaded it. Provide the corresponding VLAN ID.
   b. If you select NSX-V, provide the manager and the controller information.
   c. If you select NSX-T, provide the managers' details and the cluster/virtual IP address.
9. The NSX Controllers details for the selected platform are displayed. Click **Next**.
10. From Release 3.9 onwards, you are allowed to select the pNICs used for the NSX-T Virtual Distributed Switch (N-VDS). You can select the pNICS based on the same network speed and the available status. Heterogeneous selection on the pNIC pair is not permitted.. All the available pNICs are displayed in the list. Select any two of them from the list and click **Next**.

    Note If there are less than two pNICS available, an error message is shown and **Next** is disabled.
11. Enter the license keys for NSX for vSphere and VMware vSAN. Click **Next**.
12. Review the details. Click **Finish**. The process of adding the VxRail cluster is triggered.

Expand the VxRail Cluster

Once a cluster has been added to a workload domain, you can expand it further by adding hosts.

The process of expanding the VxRail cluster for a workload domain involves three steps:

1. Image the new node.
2. Discover and add new node to the cluster using the **Add VxRail hosts** option in the vCenter plug in of VxRail.
   a. Go to the vCenter for VxRail.
   b. Click on that particular cluster. You can view the configuration.
   c. Under the **Configure** tab, the Add VxRail Hosts page shows the list of the discovered hosts.
   d. Click **Add** to trigger the addition of hosts to that particular cluster.
   e. Provide the user credentials. Click **Next**.
   f. Review the Host IP address details. Click **Next**.
   g. Provide the ESXi credentials. Click **Next**.
   h. (Optional Step) If required, you can select the **Maintenance mode**. Click **Next**.
   i. Click **Validate**.
3. Add the host to the VMware Cloud Foundation domain cluster. The next section provides more details about this task.

Add the VxRail Hosts to the Cluster in VMware Cloud Foundation

Once the hosts have been added to the VxRail cluster, you can add them to the cluster in VMware Cloud Foundation.

Procedure

1. On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

   The **Workload Domains** page displays information for all workload domains.

2. In the workload domains table, click the name of the workload domain that you want to expand.

   The detail page for the selected workload domain appears.

3. Click the **Clusters** tab.
4. Click the name of the cluster where you want to add a host.
5. The details page for that particular cluster appears.
6. Click **Actions > Add VxRail Hosts**.

   The **Add Host To VxRail Cluster** wizard appears.

7 You can see the list of the discovered hosts. Click **Add**.

8 Under the **tasks** tab in the **cluster details** page, you can see the status of the add host task. Wait until the action is complete before performing additional workload domain tasks.

Cluster Spanning for VMware Cloud Foundation on VxRail

From Release 3.9 onwards, VMware Cloud Foundation on VxRail allows you to expand a cluster with the hosts residing in multiple L2 network domains (different VLANs and subnets).

For the management domain or the first cluster in the workload domain, the management VLAN should be stretched across all the L2 domains where the hosts of the cluster reside.

For information on expanding a cluster that needs cluster spanning, follow the steps in the VxRail documentation. For information on adding hosts, see Add the VxRail Hosts to the Cluster in VMware Cloud Foundation.

Reduce a Workload Domain

You can reduce a workload domain by removing a host from a cluster in the workload domain or by deleting a cluster.

The two tasks involved in removing a host are:

- Remove the host from the VMware Cloud Foundation domain.
- Remove the host from the vCenter plug-in of VxRail.

Remove a Host from a Cluster in a Workload Domain

You can remove a host from a cluster in a workload domain through the **Workload Domains** page in the SDDC Manager Dashboard.

When a host is removed, the vSAN members are reduced. Ensure that you have enough hosts remaining to facilitate the configured vSAN availability. Failure to do so might result in the datastore being marked as read-only or in data loss.

Procedure

1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

2 In the workload domains table, click the name of the workload domain that you want to modify.

The detail page for the selected workload domain appears.

3 Click the **Clusters** tab.

4 Click the name of the cluster from which you want to remove a host.

5 Click the **Hosts** tab.

6 Select the host to remove and click **Remove Selected Hosts**.

7 Click **Remove** to confirm the action.

The details page for the cluster appears with a message indicating that the host is being removed. When the removal process is complete, the host is removed from the hosts table.

Remove Host using the vCenter for VxRail

Once you remove the VxRail host in VMware Cloud Foundation domain, it is moved to the Maintenance mode in the vCenter for VxRail.

Perform the following tasks in the vCenter for VxRail:

Procedure

1 Right click the particular cluster and select **VxRail** -> **Remove VxRail Host**.

2 Provide the credentials. Click **Apply** to remove the host.

Delete a VxRail Cluster

You can delete a VxRail cluster from a workload domain.

You cannot delete the last cluster in a workload domain. Instead, delete the workload domain.

Prerequisites

Migrate or backup the VMs and data on the data store associated with the cluster to another location.

Procedure

1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

2 Click the name of the workload domain that contains the cluster you want to delete.

3 Click the **Clusters** tab to view the clusters in the workload domain.

4 Hover your mouse in the cluster row you want to delete.

5 Click the three dots next to the cluster name and click **Delete VxRail Cluster**.

6 Click **Delete Cluster** to confirm that you want to delete the cluster.

The details page for the workload domain appears with a message indicating that the cluster is being deleted. When the removal process is complete, the cluster is removed from the clusters table.

Delete a Workload Domain

When you delete a workload domain, the clusters within that workload domain are deleted. Note that a workload domain that includes an Edge cluster deployed cannot be deleted.

Caution Deleting a workload domain is an irreversible operation. All clusters and VMs within the workload domain are deleted and the underlying datastores are destroyed.

It can take up to 20 minutes for a workload domain to be deleted. During this process, you cannot perform any operations on workload domains.

Prerequisites

- Back up the data on the workload domain.
- Migrate the VMs that you want to keep to another workload domain.
- Shut down all VMs other than the VxRail Manager VM.

Procedure

1. On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

   The Workload Domains page displays information for all workload domains.

2. Hover your mouse in the workload domain row that where you want to delete.

   When you select the workload domain, three vertical dots appear next to the name.

3. Click the dots and choose **Delete Domain**.

   A confirmation window appears with details about the impact of deleting the workload domain, including how many hosts will be returned to the free pool.

4. Click **Delete Domain** to proceed.

   The details page for the workload domain appears with a message indicating that the workload domain is being deleted. When the removal process is complete, the workload domain is removed from the domains table.

Deploy a Workload Domain or Cluster at a Remote Location

12

With Cloud Foundation Remote Clusters, you can deploy a workload domain or cluster at a remote site through SDDC Manager. The remote workload domains and cluster are managed by the Cloud Foundation instance at the central site. You can perform a full-stack life cycle management for the remote sites from the central SDDC Manager. See Deploy a Workload Domain or Cluster at a Remote Location in *VMware Cloud Foundation Operations and Administration Guide*.

Managing Multiple Cloud Foundation Instances

13

With VMware Cloud Foundation on Dell VMC VxRail, you can manage multiple Cloud Foundation instances from a single console.

For information on using this feature, see Managing Multiple Cloud Foundation Instances in the *VMware Cloud Foundation Operations and Administration Guide*.

Stretching Clusters

14

The vSAN stretched cluster extends the cluster from one data site to two sites for high availability and load balancing. You can stretch the cluster on the management domain as well as the workload domain using the Supportability and Serviceability Utility (SoS).

You may want to stretch a cluster for the following reasons.

- Planned maintenance

  You can perform a planned maintenance on an availability zone without any downtime and then migrate the applications after the maintenance is completed.

- Automated recovery

  Stretching a cluster automatically initiates VM restart and recovery, and has a low recovery time objective for the majority of unplanned failures.

- Disaster avoidance

  With a stretched cluster, you can prevent service outages before an impending disaster such as a hurricane or rising flood levels.

Prerequisites

You have to manually validate the parameters such as the FQDN of the hosts, status of the stretched cluster, and so on, before passing it to SoS.

Availability Zones

An availability zone is a collection of infrastructure components. Each availability zone is isolated from other availability zones to prevent the propagation of failure or outage across the data center. An availability zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. Each zone should have independent power, cooling, network, and security. Additionally, these zones should be physically separate so that even uncommon disasters affect only one zone.

The physical distance between availability zones is short enough to offer low, single-digit latency (less than 5 ms) and large bandwidth (10 Gbps) between the zones. Hence, availability zones can either be two distinct data centers in a metro distance, or two safety or fire sectors (data halls) in the same large-scale data center.

The recommended minimum number of hosts in each availability zone is 4 hosts and the maximum is 15 hosts. If you are expanding a cluster, you must add hosts in pairs. Each host in the pair must have the same CPU, memory, and storage.

A region is a Cloud Foundation instance.

Note Cloud Foundation supports stretching a cluster across two availability zone within a region.

Supportability and Serviceability Utility

The SoS utility is a command-line Python tool that can be used for the following:

- Run health checks.
- On-demand vSAN partition cleanup.
- Collect logs for Cloud Foundation components.
- On-demand host cleanup.

To run the SoS utility, SSH in to the SDDC Manager VM using the **vcf** administrative user account, enter **su** to switch to the root user, and navigate to the `/opt/vmware/sddc-support` directory and type `./sos` followed by the options required for your desired operation.

```
./sos --option-1 --option-2 --option-3 ... --option-n
```

To list the available command options, use the `--help` long option or the `-h` short option.

```
./sos --help
./sos -h
```

Note You can specify some options in the conventional GNU/POSIX syntax, using -- for the long option and - for the short option.

Prerequisites for Stretching a Cluster

This section lists the prerequisites for preparing the cluster for stretching.

- Download the Deployment for Multiple Availability Zones document and read it to understand the requirements.
- Ensure that you have a vSAN Enterprise license, which is required for stretching a cluster.
- For L2 networks in NSX-V, the management VLAN, vSAN VLAN, and vMotion VLAN between the two availability zones have to be stretched.
- For L3 networks in NSX-V, the management VLAN has to be stretched.
- All VMs on an external network must be on a virtual wire. If they are on a VLAN, that VLAN must be stretched as well.
- Each availability zone must have its own vMotion, vSAN, and VXLAN networks.

- Each stretched cluster requires a vSAN witness host in a third party location. The maximum RTT on the witness is 200ms.
- The witness host should be running the same version of ESXi as the ESXi hosts in the stretched cluster. To upgrade ESXi on the witness host, use vSphere Update Manager.
- If you are stretching a cluster in a VI workload domain, you must stretch the management domain cluster first. vCenter Servers for all workload domains are in the management domain. Hence, you must protect the management domain to ensure that you can access and manage the workload domains.
- Ensure that you have enough hosts such that there is an equal number of hosts on each availability zone. This is to ensure that there are sufficient resources in case an availability zone goes down completely.
- The TCP port and the UDP ports needs to be open for witness traffic between the witness host and the vSAN cluster data nodes. See KB article 52959.

Note You cannot deploy a PKS solution on a stretched cluster.

Stretch A Cluster for NSX-V in VMware Cloud Foundation 3.9

VMware Cloud Foundation supports stretching a cluster across two availability zone within a region.

For more information on the availability zones, see Availability Zones.

Note

- Dell EMC VxRail does not support stretching clusters over the L3 networks. As a result, the cluster needs to be stretched over L2.
- Command such as `--show-free-hosts operations` is not applicable for the Dell EMC VxRail environment. If you run these commands, `'Operation is not applicable for this platform!!'` warning is thrown.
- The **Import VxRail Host** to cluster operation is disabled for a stretched cluster.

The details for the Cloud Foundation networks for Layer 2 are as follows:

Network Name	Connectivity to AZ2	Minimum MTU	Maximum MTU
vSAN	L2	1500	9000
vMotion	L2	1500	9000
VXLAN (VTEP)	L2	1600	9000
Management	L2	1500	9000

Network Name	Connectivity to AZ2	Minimum MTU	Maximum MTU
Witness Management	L3	1500	9000
Witness vSAN	L3	1500	9000

To stretch a cluster for VMware Cloud Foundation on Dell EMC VxRail, perform the following steps:

Procedure

1 Using SSH, log in to the SDDC Manager VM with the user name `vcf` and the password you specified in the deployment parameter sheet.

2 Execute the below command to prepare the cluster to be stretched. The command creates affinity rules for the VMs to run on the availability zone 1:

```
curl -X POST -H “Content-Type:application/json” http://127.0.0.1/domainmanager/vxrail/
stretch-clusters/<domain-name>/<cluster-name>?action=prepare
```

Once the workflow is triggered, track the task status in the SDDC Manager UI.

3 Power on the hosts that are rack mounted in Availability Zone 2 (or Region B) and log in to VxRail Manager. Perform the cluster expansion by adding the hosts to the cluster that is already created.

4 You have to deploy the witness in a different site. Add the witness host or the appliance to the management or the workload domain vCenter. Follow the steps listed below as described in Deploying a VSAN Witness Appliance to add a vSAN witness.

a Deploy and Configure the vSAN Witness Host in Region B.

b Add Static Routes for both Availability Zones and the vSAN Witness Host.

If the default gateway in the vSAN network provided for the network pool does not provide routing between the two availability zones and the witness host, perform all the steps in this procedure.

c Check the connectivity between the vSAN VM kernel adapters in the two availability zones and the witness host by following the instructions in KB article 1003728. Resolve errors, if any, before proceeding to the next step.

d Configure vSAN Stretched Cluster for the Management Cluster in Region A.

In step 5 of the section "Update the vSphere High Availability Settings of the Management Cluster in Region A", set **Host failures cluster tolerates** to the number of hosts in AZ1.

Note Skip the *Update Host Profiles* section to capture the vSAN stretched cluster configuration.

5 Run the following command to stretch the cluster:

```
/opt/vmware/sddc-support/sos --stretch-vsan --sc-domain <SDDC-valid-domain-name> --sc-cluster <valid cluster name which is a part of the domain to be stretched> --sc-hosts <valid host names> --witness-host-fqdn <witness host/appliance IP or fqdn> --witness-vsan-ip <witness vsan IP address> --witness-vsan-cidr <witness-vsan-network-IP-address-with-mask>
```

Enter the inputs for the following:

- `esxi host passwords`
- `vsan gateway IP`
- `vSAN CIDR`

The hosts in the Availability Zone 1 (AZ1) and the Availability Zone 2 (AZ2) are a part of the single cluster. Ensure that you name the hosts in AZ1 and AZ2 differently so that it is easy to identify. Power on the hosts that are rack mounted in AZ2 (or Region B) and log-in to VxRail Manager. As the hosts are in stretched L2 network, the VxRail Manager in AZ1 discovers the hosts automatically. Perform the cluster expansion by adding the hosts to the cluster already created in the previous steps.

6 Once the workflow is triggered, the task is tracked in the SDDC Manager UI.

7 Monitor the progress of the AZ2 hosts being added to the cluster.

a On the SDDC Manager Dashboard, click **View All Tasks**.

b Refresh the window to monitor the status.

8 Validate that stretched cluster operations are working correctly by logging in to the vSphere Web Client.

   a Verify the vSAN Health page.

      1 On the home page, click **Host and Clusters** and then select the stretched cluster (SDDC-Cluster1 in our example).

      2 Click **Monitor > vSAN > Health**.

      3 Click Retest.

      4 Fix errors, if any.

   b Verify the vSAN Storage Policy page.

      1 On the home page, click **Policies and Profiles > VM Storage Policies > vSAN Default Storage Policies** .

      2 Select the policy associated with the vCenter Server for the stretched cluster.

      3 Click **Monitor > VMs and Virtual Disks**.

      4 Click **Refresh**.

      5 Click **Trigger VM storage policy compliance check**

      6 Verify the **Compliance Status** column for each VM component.

      7 Fix errors, if any.

9 You can expand the stretched cluster.

   a Use the VxRail vCenter plugin to add the additional hosts in AZ1 or AZ2 to the cluster by performing the VxRail Manager cluster expansion work flow. Refer to the Dell EMC VxRail documentation for more details.

   b Log in to SDDC Manager and run the SoS tool to trigger the workflow to import the newly added hosts in the SDDC Manager inventory. Run the following SoS command to expand the stretched cluster:

   ```
   /opt/vmware/sddc-support/sos --expand-stretch-cluster --sc-domain <SDDC-valid-domain-name> --sc-cluster <valid cluster name which is a part of the domain to be stretched> --sc-hosts <valid host names> --witness-host-fqdn < witness host/appliance IP or fqdn> --witness-vsan-ip <witness-vsan-network-IP-address-with-mask> --witness-vsan-cidr <IP address with mask> --vsan-gateway-ip <host-vsan-gateway-ip-address>
   ```

   For both stretch and expand workflows, once the SoS command triggers, it prompts for passwords for the hosts given as inputs so you have to keep them ready in advance. In case of the expand workflow, you have to provide the fault domain information as an input for hosts. So keep the fault domain information ready.

   Enter the inputs for the following:

   - `esxi host passwords`
   - `fault domain for the hosts`
   - `vsan gateway IP`
   - `vSAN CIDR`

   Note

   - Ensure that you have the fault domain information (preferred fault domain information) for the hosts.
   - Ensure that the passwords are correct for each host.
   - For `--sc-hosts` *<valid host names>*, ensure that the multiple host names are separated by commas.

   In the SoS tool, provide the root credential and the fault domain to which the host to be added for each host.

   c Once the workflow is triggered, track the task status in the SDDC Manager UI.

Stretch a Cluster for NSX-V in VMware Cloud Foundation 3.9.1

VMware Cloud Foundation 3.9.1 supports stretching a cluster across two availability zone within a region for both Layer 2 and Layer 3 networks.

For more information on the availability zones, see Availability Zones.

Note

- Command such as `--show-free-hosts operations` is not applicable for the Dell EMC VxRail environment. If you run these commands, `'Operation is not applicable for this platform!!'` warning is thrown.
- The **Import VxRail Host** to cluster operation is disabled for a stretched cluster.
- For L2 networks in NSX-V, the management VLAN, vSAN VLAN, and vMotion VLAN between the two availability zones have to be stretched.
- For L3 networks in NSX-V, the management VLAN has to be stretched.

The details for the Cloud Foundation networks for Layer 2 are as follows:

Network Name	Connectivity to AZ2	Minimum MTU	Maximum MTU
vSAN	L2	1500	9000
vMotion	L2	1500	9000
VXLAN (VTEP)	L2	1600	9000
Management	L2	1500	9000
Witness Management	L3	1500	9000
Witness vSAN	L3	1500	9000

The details for the Cloud Foundation networks for Layer 3 are as follows:

Network Name	Connectivity to AZ2	Minimum MTU	Maximum MTU
vSAN	L3	1500	9000
vMotion	L3	1500	9000
VXLAN (VTEP)	L3	1600	9000
Management	L2	1500	9000
Witness Management	L3	1500	9000
Witness vSAN	L3	1500	9000

To stretch a cluster for VMware Cloud Foundation on Dell EMC VxRail, perform the following steps:

Procedure

1. Using SSH, log in to the SDDC Manager VM with the user name `vcf` and the password you specified in the deployment parameter sheet.

2 Prepare the workflow. Use the SoS commands to prepare the cluster. See *SoS Utility Options for vSAN Stretched Clusters* in SoS Utility Options.

   ```
   /opt/vmware/sddc-support/sos --prepare-stretch --sc-domain <SDDC-valid-domain-name> --sc-cluster <valid cluster name which is a part of the domain to be stretched>
   ```

   Once the workflow is triggered, track the task status in the SDDC Manager UI.

3 Power on the hosts that are rack mounted in Availability Zone 2 (or Region B) and log in to VxRail Manager. Perform the cluster expansion. For information on cluster expansion, refer the Dell EMC VxRail documentation.

4 You have to deploy the witness in a different site. Add the witness host or the appliance to the management or the workload domain vCenter. Follow the steps listed below as described in Deploying a VSAN Witness Appliance to add a vSAN witness.

   a Deploy and Configure the vSAN Witness Host in Region B.

   b Add Static Routes for both Availability Zones and the vSAN Witness Host.

      If the default gateway in the vSAN network provided for the network pool does not provide routing between the two availability zones and the witness host, perform all the steps in this procedure.

   c Check the connectivity between the vSAN VM kernel adapters in the two availability zones and the witness host by following the instructions in KB article 1003728. Resolve errors, if any, before proceeding to the next step.

   d Configure vSAN Stretched Cluster for the Management Cluster in Region A.

      In step 5 of the section "Update the vSphere High Availability Settings of the Management Cluster in Region A", set **Host failures cluster tolerates** to the number of hosts in AZ1.

      Note Skip the *Update Host Profiles* section to capture the vSAN stretched cluster configuration.

5 You can stretch the cluster on either Layer 2 (L2) networks or on Layer 3 (L3) networks.

a To stretch the cluster for Layer 2 (L2) networks, run the following command:

```
/opt/vmware/sddc-support/sos --l2-stretch --stretch-vsan --sc-domain <SDDC-valid-
domain-name> --sc-cluster <valid cluster name which is a part of the domain to be
stretched> --sc-hosts <valid host names> --witness-host-fqdn <witness host/appliance
IP or fqdn> --witness-vsan-ip <witness vsan IP address> --witness-vsan-cidr <witness-
vsan-network-IP-address-with-mask>
```

Note `--witness-host-fqdn` accepts either an IP address or an FQDN. If you deployed the witness host in Step 4 using an FQDN, enter the FQDN. If you deployed the witness host using an IP address, enter the IP address.

Enter the inputs for the following:

- `esxi host passwords`
- `vsan gateway IP`
- `vSAN CIDR`

For example:

```
"/opt/vmware/sddc-support/sos --l2-stretch --stretch-vsan --sc-domain MGMT --sc-cluster
VxRail-Virtual-SAN-Cluster-4fb50f56-953b-4acb-b9aa-de3f664126f6 --sc-hosts
dr27b-011.rainpole.local --witness-host-fqdn 172.27.160.106 --witness-vsan-ip 172.27.164.106
--witness-vsan-cidr 172.27.164.0/22
Welcome to Supportability and Serviceability(SoS) utility!
Logs : /var/log/vmware/vcf/sddc-support/stretchCluster-2019-10-22-07-24-05-36774
Stretch Cluster operation log : /var/log/vmware/vcf/sddc-support/
stretchCluster-2019-10-22-07-24-05-36774/sos.log
Starting vSAN stretched cluster operations..
Initiating L2 vSAN stretch operation
[**IMPORTANT**] Please make sure passwords are correct for each esxi host!!
* Please provide root user password for host dr27b-011.rainpole.local :
* Please confirm root user password for host dr27b-011.rainpole.local :
Please enter vSAN Gateway IP? (ex: 172.18.93.1): 172.18.94.1
Please enter vSAN CIDR? (ex: 172.18.93.0/24): 172.18.94.0/24
Api Response:{"id":"6b168b3f-2df8-4c84-
a431-0a8c7a9ecc6d","link":null,"taskId":"6b168b3f-2df8-4c84-
a431-0a8c7a9ecc6d","resourceId":"330c0c4a-8b4c-475d-8873-9d8db4329754","resourceType":"ESXI","
state":"IN_PROGRESS","description":"Extends VxRail vSAN cluster from a single data site to
two sites","errors":null,"timestamp":1571729075072}
Workflow triggered, please track the task status in SDDC Manager UI"
```

b To stretch the cluster for Layer 3 (L3) networks, run the following command:

```
/opt/vmware/sddc-support/sos --l3-stretch --stretch-vsan --sc-domain <SDDC-valid-
domain-name> --sc-cluster <valid cluster name which is a part of the domain to be
stretched> --sc-hosts <valid host names> --witness-host-fqdn <witness host/appliance
IP or fqdn> --witness-vsan-ip <witness vsan IP address> --witness-vsan-cidr <witness-
vsan-network-IP-address-with-mask>
```

Enter the inputs for the following:

- `esxi host passwords`
- `vsan gateway IP for the preferred(primary) and non-preferred(secondary) site`
- `vSAN CIDR for the preferred(primary) and non-preferred(secondary) site`
- `nsx vlan id`

For example:

```
root@wdc1sddc-1 [ /home/.feature ]# /opt/vmware/sddc-support/sos --l3-stretch --stretch-vsan
--sc-domain wld-1 --sc-cluster VxRail-Virtual-SAN-Cluster --sc-hosts
wdc1-005.vxrail.local,wdc1-006.vxrail.local --witness-host-fqdn 172.16.10.125 --witness-vsan-
ip 172.16.11.222 --witness-vsan-cidr 172.16.11.0/24
Welcome to Supportability and Serviceability(SoS) utility!
Logs : /var/log/vmware/vcf/sddc-support/stretchCluster-2019-11-06-12-18-02-65677
Stretch Cluster operation log : /var/log/vmware/vcf/sddc-support/
stretchCluster-2019-11-06-12-18-02-65677/sos.log
Starting vSAN stretched cluster operations..
Initiating L3 vSAN stretch operation
[**IMPORTANT**] Please make sure passwords are correct for each esxi host!!
* Please provide root user password for host wdc1-005.vxrail.local :
* Please confirm root user password for host wdc1-005.vxrail.local :
* Please provide root user password for host wdc1-006.vxrail.local :
* Please confirm root user password for host wdc1-006.vxrail.local :
** Please enter Preferred(Primary) site network information
Please enter vSAN Gateway IP? (ex: 172.18.93.1): 172.16.43.253
Please enter vSAN CIDR? (ex: 172.18.93.0/24): 172.16.43.0/24
** Please enter Non-Preferred(secondary) site network information
Please enter vSAN Gateway IP? (ex: 172.18.93.1): 172.16.11.253
Please enter vSAN CIDR? (ex: 172.18.93.0/24): 172.16.11.0/24
Please enter Preferred site Nsx Vlan Id? (ex: 800): 2057
Please enter Non-Preferred site Nsx Vlan Id? (ex: 1300): 2013
Api Response:{"id":"7fb4a115-e519-4eda-a3b7-5a7a928fc2da","link":null,"taskId":"7fb4a115-
e519-4eda-a3b7-5a7a928fc2da","resourceId":"741cd932-a01d-4fd8-
af2a-33c0a5b0317e","resourceType":"ESXI","state":"IN_PROGRESS","description":"Extends VxRail
vSAN cluster from a single data site to two sites","errors":null,"timestamp":1573042747188}
Workflow triggered, please track the task status in SDDC Manager UI
```

Note

- Ensure that the passwords are correct for each host.
- For `--sc-hosts` *`<valid host names>`*, ensure that the multiple host names are separated by commas.
- Ensure that the `witness host ip or fqdn` should match to how it is managed in vCenter. For example, if the witness host is managed using IP address in the vCenter Server, then the IP address should be provided and if the witness host is managed using FQDN in the vCenter Server, then FQDN should be provided.

6. Once the workflow is triggered, the task is tracked in the SDDC Manager UI.

7. Monitor the progress of the AZ2 hosts being added to the cluster.
   a. On the SDDC Manager Dashboard, click **View All Tasks**.
   b. Refresh the window to monitor the status.
8. Validate that stretched cluster operations are working correctly by logging in to the vSphere Web Client.
   a. Verify the vSAN Health page.
      1. On the home page, click **Host and Clusters** and then select the stretched cluster (SDDC-Cluster1 in our example).
      2. Click **Monitor > vSAN > Health**.
      3. Click Retest.
      4. Fix errors, if any.
   b. Verify the vSAN Storage Policy page.
      1. On the home page, click **Policies and Profiles > VM Storage Policies > vSAN Default Storage Policies** .
      2. Select the policy associated with the vCenter Server for the stretched cluster.
      3. Click **Monitor > VMs and Virtual Disks**.
      4. Click **Refresh**.
      5. Click **Trigger VM storage policy compliance check**
      6. Verify the **Compliance Status** column for each VM component.
      7. Fix errors, if any.
9. You can expand the stretched cluster either on Layer 2 (L2) networks or on Layer 3 (L3) networks.
   a. Use the VxRail vCenter plugin to add the additional hosts in AZ1 or AZ2 to the cluster by performing the VxRail Manager cluster expansion work flow. Refer to the Dell EMC VxRail documentation for more details.
   b. Log in to SDDC Manager and run the SoS tool to trigger the workflow to import the newly added hosts in the SDDC Manager inventory.

      In the SoS tool, provide the root credential and the fault domain to which the host to be added for each host.

c To expand the stretched cluster for Layer 2 (L2) networks, run the following SoS command :

```
/opt/vmware/sddc-support/sos --l2-stretch --expand-stretch-cluster --sc-domain <SDDC-valid-domain-name> --sc-cluster <valid cluster name which is a part of the domain to be stretched> --sc-hosts <valid host names> --witness-host-fqdn < witness host/appliance IP or fqdn> --witness-vsan-ip <witness-vsan-network-IP-address-with-mask> --witness-vsan-cidr <IP address with mask> --vsan-gateway-ip <host-vsan-gateway-ip-address>
```

Note `--witness-host-fqdn` accepts either an IP address or an FQDN. If you deployed the witness host in Step 4 using an FQDN, enter the FQDN. If you deployed the witness host using an IP address, enter the IP address.

For both stretch and expand workflows, once the SoS command triggers, it prompts for passwords for the hosts given as inputs so you have to keep them ready in advance. In case of the expand workflow, you have to provide the fault domain information as an input for hosts. So keep the fault domain information ready.

Enter the inputs for the following:

- `esxi host passwords`
- `fault domain for the hosts`
- `vSAN CIDR`

Note

- Ensure that you have the fault domain information (preferred fault domain information) for the hosts.
- Ensure that the passwords are correct for each host.
- For `--sc-hosts` *`<valid host names>`*, ensure that the multiple host names are separated by commas.
- Ensure that the `witness host ip or fqdn` should match to how it is managed in vCenter. For example, if the witness host is managed using IP address in the vCenter Server, then the IP address should be provided and if the witness host is managed using FQDN in the vCenter Server, then FQDN should be provided.

d To expand the stretched cluster for Layer 3 (L3) networks, run the following SoS command :

```
/opt/vmware/sddc-support/sos --l3-stretch --expand-stretch-cluster --sc-domain <SDDC-valid-domain-name> --sc-cluster <valid cluster name which is a part of the domain to be stretched> --sc-hosts <valid host names> --witness-host-fqdn < witness host/appliance IP or fqdn> --witness-vsan-ip <witness-vsan-network-IP-address-with-mask> --witness-vsan-cidr <IP address with mask> --vsan-gateway-ip <host-vsan-gateway-ip-address>
```

For both stretch and expand workflows, once the SoS command triggers, it prompts for passwords for the hosts given as inputs so you have to keep them ready in advance. In case of the expand workflow, you have to provide the fault domain information as an input for hosts. So keep the fault domain information ready.

Enter the inputs for the following:

- esxi host passwords
- vsan gateway IP for the preferred(primary) and non-preferred(secondary) site
- vSAN CIDR for the preferred(primary) and non-preferred(secondary) site
- nsx vlan id

For example:

```
root@wdc1sddc-1 [ /home/vcf ]# /opt/vmware/sddc-support/sos --l3-stretch --expand-stretch-
cluster --sc-domain wld-1 --sc-cluster VxRail-Virtual-SAN-Cluster --sc-hosts
wdc1-010.vxrail.local,wdc3-008.vxrail.local --witness-host-fqdn 172.16.10.125 --witness-vsan-
ip 172.16.11.222 --witness-vsan-cidr 172.16.11.0/24
Welcome to Supportability and Serviceability(SoS) utility!
Logs : /var/log/vmware/vcf/sddc-support/stretchCluster-2019-11-12-10-56-26-88147
Stretch Cluster operation log : /var/log/vmware/vcf/sddc-support/
stretchCluster-2019-11-12-10-56-26-88147/sos.log
Starting vSAN stretched cluster operations..
Initiating L3 expand vSAN stretch operation
[**IMPORTANT**]
        * Please make sure passwords are correct for each esxi host!!
        * Please keep fault-domain info handy for hosts!!
* Please provide root user password for host wdc1-010.vxrail.local :
* Please confirm root user password for host wdc1-010.vxrail.local :
* Please provide fault domain for host wdc1-010.vxrail.local :VxRail-Virtual-SAN-Cluster_az2-
faultdomain
* Please provide root user password for host wdc3-008.vxrail.local :
* Please confirm root user password for host wdc3-008.vxrail.local :
* Please provide fault domain for host wdc3-008.vxrail.local :VxRail-Virtual-SAN-Cluster_az1-
faultdomain
** Please enter Preferred(Primary) site network information
Please enter vSAN Gateway IP? (ex: 172.18.93.1): 172.16.43.253
Please enter vSAN CIDR? (ex: 172.18.93.0/24): 172.16.43.0/24
** Please enter Non-Preferred(secondary) site network information
Please enter vSAN Gateway IP? (ex: 172.18.93.1): 172.16.11.253
Please enter vSAN CIDR? (ex: 172.18.93.0/24): 172.16.11.0/24
```

```
Api Response:{"id":"c63358c5-b811-4394-b08e-ad4a42c06c19","link":null,"taskId":"c63358c5-
b811-4394-b08e-ad4a42c06c19","resourceId":"2271600f-
aee2-4df1-85e3-1e65adc075fa","resourceType":"ESXI","state":"IN_PROGRESS","description":"Expand
s VxRail vSAN stetch cluster","errors":null,"timestamp":1573556282386}
Workflow triggered, please track the task status in SDDC Manager UI
```

Note

- Ensure that you have the fault domain information (preferred fault domain information) for the hosts.
- Ensure that the passwords are correct for each host.
- For `--sc-hosts` *<valid host names>*, ensure that the multiple host names are separated by commas.
- Ensure that the `witness host ip or fqdn` should match to how it is managed in vCenter. For example, if the witness host is managed using IP address in the vCenter Server, then the IP address should be provided and if the witness host is managed using FQDN in the vCenter Server, then FQDN should be provided.

e Once the workflow is triggered, track the task status in the SDDC Manager UI.

Stretch a Cluster for NSX-T in VMware Cloud Foundation 3.9

From Release 3.9 onwards, VMware Cloud Foundation on VxRail enable L3 awareness, but the management VLAN must still be stretched between AZs.

For more information on the availability zones, see Availability Zones.

Note

- Command such as `--show-free-hosts operations` is not applicable for the Dell EMC VxRail environment. If you run these commands, `'Operation is not applicable for this platform!!'` warning is thrown.
- The **Import VxRail Host** for cluster operation is disabled for a stretched cluster.

Details for the Cloud Foundation networks:

Network Name	Connectivity to AZ2	Minimum MTU	Maximum MTU
vSAN	L3	1500	9000
vMotion	L3	1500	9000
VXLAN (VTEP)	L3	1600	9000
Management	L2	1500	9000
Witness Management	L3	1500	9000
Witness vSAN	L3	1500	9000

To stretch a cluster for VMware Cloud Foundation on Dell EMC VxRail, perform the following steps:

Procedure

1. Using SSH, log in to the SDDC Manager VM with the user name `vcf` and the password you specified in the deployment parameter sheet.

2. Execute the below command to prepare the cluster to be stretched. The command creates affinity rules for the VMs to run on the availability zone 1:

   ```
   curl -X POST -H “Content-Type:application/json” http://127.0.0.1/domainmanager/vxrail/stretch-clusters/<domain-name>/<cluster-name>?action=prepare
   ```

   Once the workflow is triggered, track the task status in the SDDC Manager UI.

3. You have to deploy the witness in a different site. Add the witness host or the appliance to the management or the workload domain vCenter. Follow the steps listed below as described in Deploying a VSAN Witness Appliance to add a vSAN witness.

   a Deploy and Configure the vSAN Witness Host in Region B.

   b Add Static Routes for both Availability Zones and the vSAN Witness Host.

   If the default gateway in the vSAN network provided for the network pool does not provide routing between the two availability zones and the witness host, perform all the steps in this procedure.

   c Check the connectivity between the vSAN VM kernel adapters in the two availability zones and the witness host by following the instructions in KB article 1003728. Resolve errors, if any, before proceeding to the next step.

   d " Configure vSAN Stretched Cluster for the Management Cluster in Region A"

   In step 5 of the section "Update the vSphere High Availability Settings of the Management Cluster in Region A", set **Host failures cluster tolerates** to the number of hosts in AZ1.

   Note Skip the section Update Host Profiles to Capture the vSAN Stretched Cluster Configuration.

4. Expand cluster with Availability Zone 2 hosts in vCenter. Reimage and power on the hosts that are rack mounted in Availability Zone 2 (or Region B). See Cluster Spanning for VMware Cloud Foundation on VxRail for more information on expanding clusters that need cluster spanning.

5. To add a host, in SDDC Manager, select the NSX-T workload domain cluster and perform the **Add Host** operation. This binds the Availability Zone 2 hosts to NSX-T Manager and vRealize Log Insight. Once this operation is completed, the Availability Zone 2 hosts are available in SDDC Manager inventory.

6. Configure the static routes on the ESXi hosts. On each of the ESXi hosts configure the static route to reach the Witness host/appliance. To configure the static route on ESXi host, see https://kb.vmware.com/s/article/2001426.

7 Configure the vSphere high availability options as follows:

   a Select the cluster to be stretched in vCenter.

   b Navigate to **Configure > Services > Availability**.

   c Click **Edit** and navigate to the **Admission Control** settings.

   d In **Admission Control** settings, select **Reserved Failover CPU capacity** to **50% CPU** and **Reserved Failover Memory Capacity** to **50% Memory**.

   e Navigate to the **Advanced Options** tab and click **Add**.

   f In the new row, set the **Option** to `das.isolationaddress0` and **Value** to the vSAN gateway of Availability Zone 1. Click **Add** again and in the second new row set the **Option** to `das.isolationaddress1` and**Value** to the vSAN gateway of Availability Zone 2.

8 To stretch the cluster in vCenter, manually select the hosts that will be part of fault domain 1 and fault domain 2 and also providing the witness host information.

9 At this point, the cluster is stretched. However, as the stretch operation was performed in vCenter and not by using SDDC Manager, the cluster data in SDDC Manager is not in sync with vCenter. For more information, see Manually Configure vSAN Stretched Cluster.

10 Once the workflow is triggered, the task is tracked in the SDDC Manager UI.

11 Monitor the progress of the AZ2 hosts being added to the cluster.

   a On the SDDC Manager Dashboard, click **View All Tasks**.

   b Refresh the window to monitor the status.

12 Validate that stretched cluster operations are working correctly by logging in to the vSphere Web Client.

   a Verify the vSAN Health page.

      1 On the home page, click **Host and Clusters** and then select the stretched cluster (SDDC-Cluster1 in our example).

      2 Click **Monitor > vSAN > Health**.

      3 Click **Retest**.

      4 Fix errors, if any.

   b Verify the vSAN Storage Policy page.

      1 On the home page, click **Policies and Profiles > VM Storage Policies > vSAN Default Storage Policies** .

      2 Select the policy associated with the vCenter Server for the stretched cluster.

      3 Click **Monitor > VMs and Virtual Disks**.

      4 Click **Refresh**.

      5 Click **Trigger VM storage policy compliance check**

      6 Check the **Compliance Status** column for each VM component.

      7 Fix errors, if any.

13 You can expand the stretched cluster. After the cluster is stretched, you can add more VxRail nodes to Availability Zone 1 or Availability Zone 2 as follows:

   a In the vCenter-VxRail plugin, add the new VxRail hosts to the cluster. Ensure that the newly added hosts are in the maintenance mode.

   b In vCenter, if the new VxRail host is a part of Availability Zone 1, add the VxRail host to the host group.

      1 Select the cluster.

      2 Navigate to the Configure tab.

      3 Select **Configuration > VM/Host Groups**.

      4 Select the host group corresponding to Availability Zone 1 and add the newly added VxRail hosts.

   c In vCenter, add the new VxRail host to the fault domain for the Availability Zone.

      1 Select the cluster.

      2 Navigate to the **Configure** tab.

      3 Select **vSAN > Fault Domains**.

      4 Add the newly included VxRail hosts to the respective fault domains.

   d Add the new VxRail node to SDDC Manager.

14 You can shrink the stretched cluster by removing the nodes from the cluster.

 a Switch the VxRail hosts that are part of Availability Zone 2 into the maintenance mode.

 b Remove the hosts that are part of Availability Zone 2 from the host group.

 c Remove the host group created for Availability Zone 2.

 d Remove every removed host that is part of Availability Zone 2 from SDDC Manager.

 e Remove every removed host that is part of Availability Zone 2 from vCenter VxRail UI.

 f Log in to each ESXi server that is a part of Availability Zone 1 and remove the static routes to reach Witness vSAN host/appliance.

 See the following knowledge base article to configure the static route on the ESXi host:

 https://kb.vmware.com/s/article/2001426

 g In vCenter, remove the vSphere Availability settings that were made while stretching the cluster.

Stretch a Cluster for NSX-T in VMware Cloud Foundation 3.9.1

From Release 3.9 onwards, VMware Cloud Foundation on VxRail enable L3 awareness, but the management VLAN must still be stretched between AZs.

For more information on the availability zones, see Availability Zones.

Note

- Command such as `--show-free-hosts operations` is not applicable for the Dell EMC VxRail environment. If you run these commands, `'Operation is not applicable for this platform!!'` warning is thrown.
- The **Import VxRail Host** for cluster operation is disabled for a stretched cluster.

Details for the Cloud Foundation networks:

Network Name	Connectivity to AZ2	Minimum MTU	Maximum MTU
vSAN	L3	1500	9000
vMotion	L3	1500	9000
Tunnel End Points (TEP)	L3	1600	9000
Management	L2	1500	9000
Witness Management	L3	1500	9000
Witness vSAN	L3	1500	9000

To stretch a cluster for VMware Cloud Foundation on Dell EMC VxRail, perform the following steps:

Procedure

1 Using SSH, log in to the SDDC Manager VM with the user name `vcf` and the password you specified in the deployment parameter sheet.

2 Prepare the workflow. Use the SoS commands to prepare the cluster. See *SoS Utility Options for vSAN Stretched Clusters* in SoS Utility Options.

```
/opt/vmware/sddc-support/sos --prepare-stretch --sc-domain <SDDC-valid-domain-name> --sc-cluster <valid cluster name which is a part of the domain to be stretched>
```

Once the workflow is triggered, track the task status in the SDDC Manager UI.

3 You have to deploy the witness in a different site. Add the witness host or the appliance to the management or the workload domain vCenter. Follow the steps listed below as described in Deploying a VSAN Witness Appliance to add a vSAN witness.

a Deploy and Configure the vSAN Witness Host in Region B.

b Add Static Routes for both Availability Zones and the vSAN Witness Host.

If the default gateway in the vSAN network provided for the network pool does not provide routing between the two availability zones and the witness host, perform all the steps in this procedure.

c Check the connectivity between the vSAN VM kernel adapters in the two availability zones and the witness host by following the instructions in KB article 1003728. Resolve errors, if any, before proceeding to the next step.

d " Configure vSAN Stretched Cluster for the Management Cluster in Region A"

In step 5 of the section "Update the vSphere High Availability Settings of the Management Cluster in Region A", set **Host failures cluster tolerates** to the number of hosts in AZ1.

Note Skip the section Update Host Profiles to Capture the vSAN Stretched Cluster Configuration.

4 Run the following command to stretch the cluster for Layer 3 (L3) networks:

```
/opt/vmware/sddc-support/sos --l3-stretch --stretch-vsan --sc-domain <SDDC-valid-domain-name> --sc-cluster <valid cluster name which is a part of the domain to be stretched> --sc-hosts <valid host names> --witness-host-fqdn <witness host/appliance IP or fqdn> --witness-vsan-ip <witness vsan IP address> --witness-vsan-cidr <witness-vsan-network-IP-address-with-mask>
```

Enter the inputs for the following:

- `esxi host passwords`
- `vsan gateway IP for the preferred(primary) and non-preferred(secondary) site`
- `vSAN CIDR for the preferred(primary) and non-preferred(secondary) site`
- `nsx vlan id`

For example:

```
root@wdc1sddc-1 [ /home/vcf ]# /opt/vmware/sddc-support/sos --l3-stretch --expand-stretch-cluster
--sc-domain wld-1 --sc-cluster VxRail-Virtual-SAN-Cluster --sc-hosts
wdc1-010.vxrail.local,wdc3-008.vxrail.local --witness-host-fqdn 172.16.10.125 --witness-vsan-ip
172.16.11.222 --witness-vsan-cidr 172.16.11.0/24
Welcome to Supportability and Serviceability(SoS) utility!
Logs : /var/log/vmware/vcf/sddc-support/stretchCluster-2019-11-12-10-56-26-88147
Stretch Cluster operation log : /var/log/vmware/vcf/sddc-support/
stretchCluster-2019-11-12-10-56-26-88147/sos.log
Starting vSAN stretched cluster operations..
Initiating L3 expand vSAN stretch operation
[**IMPORTANT**]
        * Please make sure passwords are correct for each esxi host!!
        * Please keep fault-domain info handy for hosts!!
* Please provide root user password for host wdc1-010.vxrail.local :
* Please confirm root user password for host wdc1-010.vxrail.local :
* Please provide fault domain for host wdc1-010.vxrail.local :VxRail-Virtual-SAN-Cluster_az2-
faultdomain
* Please provide root user password for host wdc3-008.vxrail.local :
* Please confirm root user password for host wdc3-008.vxrail.local :
* Please provide fault domain for host wdc3-008.vxrail.local :VxRail-Virtual-SAN-Cluster_az1-
faultdomain
** Please enter Preferred(Primary) site network information
Please enter vSAN Gateway IP? (ex: 172.18.93.1): 172.16.43.253
Please enter vSAN CIDR? (ex: 172.18.93.0/24): 172.16.43.0/24
** Please enter Non-Preferred(secondary) site network information
Please enter vSAN Gateway IP? (ex: 172.18.93.1): 172.16.11.253
Please enter vSAN CIDR? (ex: 172.18.93.0/24): 172.16.11.0/24
Api Response:{"id":"c63358c5-b811-4394-b08e-ad4a42c06c19","link":null,"taskId":"c63358c5-
b811-4394-b08e-ad4a42c06c19","resourceId":"2271600f-
aee2-4df1-85e3-1e65adc075fa","resourceType":"ESXI","state":"IN_PROGRESS","description":"Expands
VxRail vSAN stetch cluster","errors":null,"timestamp":1573556282386}
Workflow triggered, please track the task status in SDDC Manager UI
```

5 Expand cluster with Availability Zone 2 hosts in vCenter. Reimage and power on the hosts that are rack mounted in Availability Zone 2 (or Region B). See Cluster Spanning for VMware Cloud Foundation on VxRail for more information on expanding clusters that need cluster spanning.

   a Use the VxRail vCenter plugin to add the additional hosts in AZ1 or AZ2 to the cluster by performing the VxRail Manager cluster expansion work flow. Refer to the Dell EMC VxRail documentation for more details.

   b Log in to SDDC Manager and run the SoS tool to trigger the workflow to import the newly added hosts in the SDDC Manager inventory.

   In the SoS tool, provide the root credential and the fault domain to which the host to be added for each host.

c Run the following SoS command to expand the stretched cluster for Layer 3 (L3) networks:

/opt/vmware/sddc-support/sos --l3-stretch --expand-stretch-cluster --sc-domain *<SDDC-valid-domain-name>* --sc-cluster *<valid cluster name which is a part of the domain to be stretched>* --sc-hosts *<valid host names>* --witness-host-fqdn *< witness host/appliance IP or fqdn>* --witness-vsan-ip *<witness-vsan-network-IP-address-with-mask>* --witness-vsan-cidr *<IP address with mask>* --vsan-gateway-ip *<host-vsan-gateway-ip-address>*

For both stretch and expand workflows, once the SoS command triggers, it prompts for passwords for the hosts given as inputs so you have to keep them ready in advance. In case of the expand workflow, you have to provide the fault domain information as an input for hosts. So keep the fault domain information ready.

Enter the inputs for the following:

- `esxi host passwords`
- `vsan gateway IP for the preferred(primary) and non-preferred(secondary) site`
- `vSAN CIDR for the preferred(primary) and non-preferred(secondary) site`
- `nsx vlan id`

For example:

```
root@wdc1sddc-1 [ /home/vcf ]# /opt/vmware/sddc-support/sos --l3-stretch --expand-stretch-
cluster --sc-domain wld-1 --sc-cluster VxRail-Virtual-SAN-Cluster --sc-hosts
wdc1-010.vxrail.local,wdc3-008.vxrail.local --witness-host-fqdn 172.16.10.125 --witness-vsan-
ip 172.16.11.222 --witness-vsan-cidr 172.16.11.0/24
Welcome to Supportability and Serviceability(SoS) utility!
Logs : /var/log/vmware/vcf/sddc-support/stretchCluster-2019-11-12-10-56-26-88147
Stretch Cluster operation log : /var/log/vmware/vcf/sddc-support/
stretchCluster-2019-11-12-10-56-26-88147/sos.log
Starting vSAN stretched cluster operations..
Initiating L3 expand vSAN stretch operation
[**IMPORTANT**]
        * Please make sure passwords are correct for each esxi host!!
        * Please keep fault-domain info handy for hosts!!
* Please provide root user password for host wdc1-010.vxrail.local :
* Please confirm root user password for host wdc1-010.vxrail.local :
* Please provide fault domain for host wdc1-010.vxrail.local :VxRail-Virtual-SAN-Cluster_az2-
faultdomain
* Please provide root user password for host wdc3-008.vxrail.local :
* Please confirm root user password for host wdc3-008.vxrail.local :
* Please provide fault domain for host wdc3-008.vxrail.local :VxRail-Virtual-SAN-Cluster_az1-
faultdomain
** Please enter Preferred(Primary) site network information
Please enter vSAN Gateway IP? (ex: 172.18.93.1): 172.16.43.253
Please enter vSAN CIDR? (ex: 172.18.93.0/24): 172.16.43.0/24
** Please enter Non-Preferred(secondary) site network information
Please enter vSAN Gateway IP? (ex: 172.18.93.1): 172.16.11.253
Please enter vSAN CIDR? (ex: 172.18.93.0/24): 172.16.11.0/24
```

```
Api Response:{"id":"c63358c5-b811-4394-b08e-ad4a42c06c19","link":null,"taskId":"c63358c5-
b811-4394-b08e-ad4a42c06c19","resourceId":"2271600f-
aee2-4df1-85e3-1e65adc075fa","resourceType":"ESXI","state":"IN_PROGRESS","description":"Expand
s VxRail vSAN stetch cluster","errors":null,"timestamp":1573556282386}
Workflow triggered, please track the task status in SDDC Manager UI
```

Note

- Ensure that you have the fault domain information (preferred fault domain information) for the hosts.
- Ensure that the passwords are correct for each host.
- For `--sc-hosts` *`<valid host names>`*, ensure that the multiple host names are separated by commas.

d Once the workflow is triggered, track the task status in the SDDC Manager UI.

6 Monitor the progress of the AZ2 hosts being added to the cluster.

a On the SDDC Manager Dashboard, click **View All Tasks**.

b Refresh the window to monitor the status.

7 Validate that stretched cluster operations are working correctly by logging in to the vSphere Web Client.

a Verify the vSAN Health page.

1 On the home page, click **Host and Clusters** and then select the stretched cluster (SDDC-Cluster1 in our example).

2 Click **Monitor > vSAN > Health**.

3 Click **Retest**.

4 Fix errors, if any.

b Verify the vSAN Storage Policy page.

1 On the home page, click **Policies and Profiles > VM Storage Policies > vSAN Default Storage Policies** .

2 Select the policy associated with the vCenter Server for the stretched cluster.

3 Click **Monitor > VMs and Virtual Disks**.

4 Click **Refresh**.

5 Click **Trigger VM storage policy compliance check**

6 Check the **Compliance Status** column for each VM component.

7 Fix errors, if any.

8 You can shrink the stretched cluster by removing the nodes from the cluster.

   a Switch the VxRail hosts that are part of Availability Zone 2 into the maintenance mode.

   b Remove the hosts that are part of Availability Zone 2 from the host group.

   c Remove the host group created for Availability Zone 2.

   d Remove every removed host that is part of Availability Zone 2 from SDDC Manager.

   e Remove every removed host that is part of Availability Zone 2 from vCenter VxRail UI.

   f Log in to each ESXi server that is a part of Availability Zone 1 and remove the static routes to reach Witness vSAN host/appliance.

      See the following knowledge base article to configure the static route on the ESXi host:

      https://kb.vmware.com/s/article/2001426

   g In vCenter, remove the vSphere Availability settings that were made while stretching the cluster.

Unstretch a Cluster

From VMware Cloud Foundation 3.7.2 onwards, you can unstretch a management or a workload domain cluster that is already stretched.

Procedure

1 Log in to SDDC Manager and run the SoS tool to trigger the workflow to unstretch a cluster. See *SoS Utility Options for vSAN Stretched Clusters* in SoS Utility Options.

2 Run the following SoS command to unstretch a cluster.

```
/opt/vmware/sddc-support/sos --unstretch-vsan --sc-domain <domain-name> --sc-cluster
<cluster-name>
```

   Track the task status of this workflow in the SDDC Manager UI.

Lifecycle Management

15

Lifecycle Management (LCM) enables you to perform automated updates on Cloud Foundation services (SDDC Manager and internal services), VMware software (NSX for vSphere, vCenter Server, ESXi, and vRSCLM), and Dell EMC VxRail in your environment. You can download the update bundles and apply them manually or schedule them within your maintenance window allowing for flexibility in your application.

For more information on VMware Cloud Foundation Lifecycle Management, see Patching and Upgrading Cloud Foundation. For information on upgrading VMware Cloud Foundation, see *VMware Cloud Foundation Upgrade Guide.*.

The LCM bundles that are available are:

- VxRail Partner Bundle: You can download the Dell EMC VxRail partner bundle to update the VxRail appliance.
- Patch Update Bundle: A patch update bundle contains bits to update the appropriate Cloud Foundation software components in your management domain or VI workload domain. In most cases, a patch update bundle must be applied to the management domain before it can be applied to VI workload domains.
- Cumulative Update Bundle: With a cumulative update bundle, you can directly update the appropriate software in your workload domain to the version contained in the cumulative bundle rather than applying sequential updates to reach the target version.
- Install Bundle: If you have updated the management domain in your environment, you can download an install bundle with updated software bits for VI workload domains and vRealize suite components.

This chapter includes the following topics:

Online Bundle Download

You must log in to My VMware account before downloading a bundle.

You must be logged in to My VMware account to download update bundles.

1 In the SDDC dashboard, click **Administration > Repository Settings** .

2 Enter the My VMware credentials as well as the Dell EMC credentials. Click **Authenticate**. You have to log into both My VMware and Dell EMC to update all the VMware Cloud Foundation and the VxRail components.

3 To access the bundles, you have two options:

   - In the SDDC Manager Dashboard, navigate to the **Bundles** page. This page shows the available update bundles for the components.

     1 Click **Repository > Bundles**.

   - In the SDDC Manager Dashboard, navigate to the **Workload Domain** page.

     1 Click **Inventory > Workload Domains**.

     2 Click the name of a workload domain and then click the **Updates/Patches** tab.

        The number next to the Updates/Patches tab indicates the available updates.

        The **Available Updates** section displays all updates applicable to this workload domain.

        Also, you can view the current versions running such as the VxRail current version.

4 To view the metadata details for an update bundle, click **View Details**. The bundle severity and detailed information about each component included in the bundle is displayed. If a bundle is a cumulative bundle, this information is displayed as well. The bundle severity levels are described in the table below.

Severity Value	Description
Critical	A problem may severely impact your production systems (including the loss of production data). Such impacts could be system down or HA not functioning.
Important	A problem may affect functionality, or cause a system to function in a severely reduced capacity. The situation causes significant impact to portions of the business operations and productivity. The system is exposed to potential loss or interruption of services. A change to support hardware enablement (for example, a driver update), or a new feature for an important product capability.
Moderate	A problem may affect partial non-critical functionality loss. This may be a minor issue with limited loss, no loss of functionality, or impact to the client's operations and issues in which there is an easy circumvention or avoidance by the end user. This includes documentation errors.
Low	A problem which has low or no impact to a product's functionality or a client's operations. There is no impact on quality, performance, or functionality of the product.

5 Click **Schedule Download**. Select the date and time for the bundle download and click **Schedule**.

Use a Proxy Server to Download Upgrade Bundles

If you do not have internet access, you can use a proxy server to download the LCM bundles. LCM only supports proxy servers that do not require authentication.

Procedure

1. Using SSH, log in to the SDDC Manager VM with the user name `vcf` and password you specified in the deployment parameter sheet.
2. Type `su` to switch to the root account.
3. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
4. Add the following lines to the end of the file:

   ```
   lcm.depot.adapter.proxyEnabled=true
   lcm.depot.adapter.proxyHost=proxy IP address
   lcm.depot.adapter.proxyPort=proxy port
   ```

5. Save and close the file.
6. Restart the LCM server by typing the following command in the console window:

   ```
   systemctl restart lcm
   ```

7. Wait for 5 minutes and then download the bundles.

Update your Environment

The SDDC Manager is the first component to be upgraded on the management domain.

Components must be updated in the following order:

1. Cloud Foundation services
2. NSX-T
3. vCenter Server
4. VxRail
5. ESXi

   If you have multiple clusters in the management domain or in a workload domain, you can upgrade VxRail at a cluster level.

Procedure

1. On the SDDC Manager Dashboard, click > **Inventory**.
2. Click a domain and then click the **Updates/Patches** tab.

3 Click **Precheck** to validate that the appliance is ready to be updated.

Click **View Status** to see the update status for each component and the tests performed. Expand a test by clicking the arrow next to it to see further details.

If any of the tests fail, fix the issue and click **Retry Precheck**.

The precheck results are displayed below the Precheck button. Ensure that the precheck results are green before proceeding. A failed precheck may cause the update to fail.

4 Do either of the following:

a Click **Update Now**.

b Click **Schedule Download**. Select the date and time for the bundle download and click Schedule.

To view the update status, click **Update History** > **Actions** >**View Update Status**. After the bundle is downloaded, the **Schedule Update** button is displayed. Click **View Details** to see the version changes for each component that the bundle will apply.

5 After the update is completed successfully, log out of the SDDC Manager Dashboard and log back in.

Support for the Manual Download of Bundles for VMware Cloud Foundation

LCM polls the VMware depot to access the update bundles. If you do not have internet connectivity in your Cloud Foundation system, you can use the Bundle Transfer utility to manually the bundles from the depot on your local computer and then upload them to SDDC Manager. The utility identifies applicable bundles based on the current software versions in your environment based on a marker file generated on the SDDC Manager VM.

Prerequisites

A Windows or Linux computer with Java 8 or later and internet connectivity for downloading the bundles.

Procedure

1 Using SSH, log in to the SDDC Manager VM with the user name `vcf` and password you specified in the deployment parameter sheet.

2 Navigate to the `/opt/vmware/vcf/lcm/lcm-tools/bin` directory.

3 Generate a marker file by running the following command.

```
./lcm-bundle-transfer-util --generateMarker
```

The marker file (markerFile) is a JSON file that contains information on the current software versions running on SDDC Manager. It also contains the bundles IDs for bundles that were downloaded before this file was generated. The markerFile.md5 contains the checksum for the markerFile.

The output contains the directory where the marker file is generated.

4 Copy the /opt/vmware/vcf/lcm/lcm-tools/bin directory, and the markerFile and markerFile.md5 files from the location displayed in the output of step 3 to a computer with internet access.

The /opt/vmware/vcf/lcm/lcm-tools/bin directory includes the bundle transfer utility required for the next step.

5 For the VxRail SKU, when you download the NSX or vCenter bundles, provide the additional argument of withCompatibilitySets which will let you download the compatibility sets along with the bundles. On the computer with internet access, run the following command.

```
./lcm-bundle-transfer-util -download "withCompatibilitySets"
            -outputDirectory ${absolute-path-output-dir}
            -depotUser ${depotUser}
            -markerFile ${absolute-path-markerFile}
            -markerMd5File ${absolute-path-markerFile.md5}
```

where

absolute-path-output-dir	Path to the directory where the bundle files are to be downloaded. This directory folder must have 777 permissions. If you do not specify the download directory, bundles are downloaded to the default directory with 777 permissions.
depotUser	User name for myVMware depot. You are prompted to enter the depot user password. If there are any special characters in the password, specify the password within single quotes.
markerFile	Absolute path to the marker file, as generated in the above step. If you do not specify the path to the marker file, all update bundles on the depot are downloaded.
markerMd5File	Absolute path to the marker MD5 checksum file, as generated in the above step.

The utility generates a delta file (deltaFileDownloaded) in the download directory based on the software versions in the marker file and the update bundles available on the depot. The applicable bundles identified in the delta file are downloaded. Download progress for each bundle is displayed.

6 When you download the VxRail bundles, provide the additional argument downloadPartnerBundle.

```
./lcm-bundle-transfer-util -download "downloadPartnerBundle"
            -outputDirectory ${absolute-path-output-dir}
            -depotUser ${vmwaredepotUser}
```

```
    -depotUserPassword ${vmwareDepotPassword}
    -pdu ${emcdepotuser}:${emcdepotpassword}
            -markerFile ${absolute-path-markerFile}
            -markerMd5File ${absolute-path-markerFile.md5}
```

7 Copy the downloaded `softwareCompatibilitySets.json` to `/nfs/vmware/vcf/nfs-mount/bundle/depot/local/softwareCompatibilitySets.json`.

8 Copy the update bundle directory from the external computer to the SDDC Manager VM.

For example:

```
scp -pr /Work/UpdateBundle vcf@SDDC_IP:/home/vcf/vCF372to38Bundle"
```

9 In the SDDC Manager VM, change the ownership and permissions of the uploaded bundle.

```
chown vcf_lcm:vcf -R /opt/vmware/vcf/vCF372to38Bundle
chmod -R 0777 /opt/vmware/vcf/vCF372to38Bundle
```

10 In the SDDC Manager VM, upload the bundle files to the internal LCM repository.

```
cd /opt/vmware/vcf/lcm/lcm-tools/bin
./lcm-bundle-transfer-util -upload "withCompatibilitySets" -bundleDirectory
${absolute-path-output-dir}
```

where *absolute-path-output-dir* is the directory where the bundle files have been be uploaded, or `/opt/vmware/vcf/vCF372to38Bundle` as shown in the previous step.

The utility uploads the bundles specified in the `deltaFileDownloaded` file. The console displays upload status for each bundle.

11 For the VxRail bundles, perform the following:

a Copy the partner bundle to `/nfs/vmware/vcf/nfs-mount/bundle/depot/local/bundles`.

b Copy `partnerBundleMetadata.json` to `/nfs/vmware/vcf/nfs-mount/bundle/depot/local/partnerBundleMetadata.json`.

c Copy `softwareCompatibilitySets.json` to `/nfs/vmware/vcf/nfs-mount/bundle/depot/local/softwareCompatibilitySets.json`

d In the SDDC Manager VM, upload the bundle files to the internal LCM repository.

```
cd /opt/vmware/vcf/lcm/lcm-tools/bin
./lcm-bundle-transfer-util -upload "uploadPartnerBundle" -bundleDirectory
${absolute-path-output-dir}
```

Update History

The **Update History** page displays all updates applied to a workload domain.

Procedure

1 In the SDDC Manager Dashboard, click **Inventory > Workload Domains**..

2 Click the name of a workload domain and then click the **Updates History** tab.

  All updates applied to this workload domain are displayed. If an update bundle was applied more than once, click **View Past Attempts** to see more information.

Image-Based Backup and Restore

16

For an image-based backup of the SDDC Manager, use a solution compatible with the VMware vSphere Storage APIs - Data Protection (formerly known as VMware vStorage APIs for Data Protection or VADP).

vSphere Storage APIs - Data Protection compatible backup software connects to the vCenter servers in the management domain to perform the backup to a remote site. In the event of failure, the backup software connects to the vCenter servers in the management domain to restore the VMs. If the management domain is lost, the vCenter servers are no longer available and must be restored first. Choosing a backup software that supports Direct Restore to a ESXi host allows restoring the vCenter servers.

For an SDDC Manager backup, connect your backup with the management domain vCenter Server. Configure the product to take non-quiesced backups of SDDC Manager. To reduce the backup time and storage cost, use incremental backups in addition to full ones.

Configuring Customer Experience Improvement Program 17

This product participates in VMware Customer Experience Improvement Program (CEIP).

The CEIP provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. As part of the CEIP, VMware collects the technical information about your organization's use of the VMware products and services regularly in association with your organization's VMware license keys. This information does not personally identify any individual. For additional information regarding the CEIP, refer to the Trust & Assurance Center at http://www.vmware.com/trustvmware/ceip.html.

You can enable or disable CEIP across all the SDDC components deployed in VMware Cloud Foundation by the following methods:

- When you log into SDDC Manager for the first time, a pop-up window appears. The **Join the VMware Customer Experience Program** option is selected by default. Deselect this option if you do not want to enable CEIP. Click **Apply**.

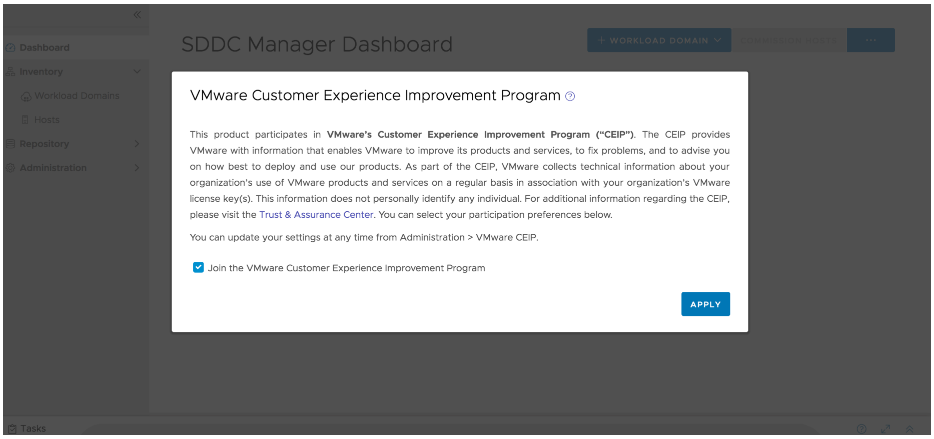

VMware Customer Experience Improvement Program

This product participates in **VMware's Customer Experience Improvement Program ("CEIP")**. The CEIP provides VMware with information that enables VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. As part of the CEIP, VMware collects technical information about your organization's use of VMware products and services on a regular basis in association with your organization's VMware license key(s). This information does not personally identify any individual. For additional information regarding the CEIP, please visit the Trust & Assurance Center. You can select your participation preferences below.

You can update your settings at any time from Administration > VMware CEIP.

☑ Join the VMware Customer Experience Improvement Program

APPLY

- You can enable or disable CEIP from the Administration tab on the SDDC Manager dashboard.

Note When you join CEIP, the **Enable VMware Customer Experience Improvement Program** task is shown in the floating task bar of the SDDC dashboard. Similarly when you leave CEIP, the **Disable VMware Customer Experience Improvement Program** task is displayed. This task bar is available on all the screens of VMware Cloud Foundation.

To enable or disable CEIP from the **Administration** tab, perform the following steps:

Procedure

1. On the SDDC Manager Dashboard, click **Administration > VMware CEIP**.
2. To enable CEIP, select the **Join the VMware Customer Experience Improve Program** option.
3. To disable CEIP, deselect the **Join the VMware Customer Experience Improve Program** option.